Heuristic energy-based cyclic peptide design

Qiyao Zhu[1], Vikram Khipple Mulligan[1], Dennis Shasha[2,*]

1 Center for Computational Biology, Flatiron Institute, New York, NY, U.S.A.
2 Courant Institute of Mathematical Sciences, New York University, New York, NY, U.S.A.

* shasha@cims.nyu.edu

Abstract

Rational computational design is crucial to the pursuit of novel drugs and therapeutic agents. Meso-scale cyclic peptides, which consist of 7-40 amino acid residues, are of particular interest due to their conformational rigidity, binding specificity, degradation resistance, and potential cell permeability. Because there are few natural cyclic peptides, *de novo* design involving non-canonical amino acids is a potentially useful goal. Here, we develop an efficient pipeline (CyclicChamp) for cyclic peptide design. After converting the cyclic constraint into an error function, we employ a variant of simulated annealing to search for low-energy peptide backbones while maintaining peptide closure. Compared to the previous random sampling approach, which was capable of sampling conformations of cyclic peptides of up to 14 residues, our method both greatly accelerates the computation speed for sampling conformations of small macrocycles (*ca.* 7 residues), and addresses the high-dimensionality challenge that large macrocycle designs often encounter. As a result, CyclicChamp makes conformational sampling tractable for 15- to 24-residue cyclic peptides, thus permitting the design of macrocycles in this size range. Microsecond-length molecular dynamics simulations on the resulting 15, 20, and 24 amino acid cyclic designs identify trajectories with kinetic stability. To test their thermodynamic stability, we perform additional replica exchange molecular dynamics simulations and generate free energy surfaces. Two 15-residue designs and one 20-residue design emerge as promising candidates, along with one viable 24-residue candidate.

Author summary

Cyclic peptides are circular chains of amino acid residues that are promising candidates for new therapeutic drugs. Current FDA approved cyclic peptide-based drugs are mostly derived from natural sources. However, recent work has enabled the computational design of new cyclic peptide drugs. Current *de novo* computational design methods can handle sizes of 7 to 13 residues without conformational constraints. As size increases, the exponentially growing conformational space makes conformational sampling intractable. The literature's prevalent approach of random sampling finds poor configurations, with the result that the success rate of finding a stable design is very low. Here, we develop an efficient search algorithm by combining tailored optimization algorithms with established energy models. Our heuristic design pipeline, CyclicChamp, produces stable cyclic peptide designs of 7, 15, 20, and 24 amino acids as validated by algorithmically-independent molecular dynamics simulations. This pipeline not only expands the structural variety for future drug development, but also paves the way for potential cyclic peptide-based enzyme design.

Introduction

Cyclic peptides are chains of fewer than 40 amino acid residues forming one or more closed loops. One common class of cyclic peptides consists of a single loop, with the N- and C-termini connected by an

amide bond. In general, mid-sized cyclic peptides stand out for their superior binding affinity and selectivity compared to small molecules [1]. Their still modest size reduces the likelihood of provoking an immune response and enhances the ability to traverse cellular barriers compared to large protein therapeutics [2]. Further, their characteristic cyclization imposes conformational constraints, leading to structures that are more rigid compared to their linear counterparts [3]. This rigidity reduces the entropic cost associated with ordering a disordered molecule on binding to its intended target, enhancing target affinity [4]; it also prevents adoption of alternative conformations in which the peptide may bind to off-target proteins, thus enhancing specificity [5]. In addition, the connection of the N- and C-termini makes cyclic peptides more resistant to protease degradation than linear peptides. The incorporation of non-canonical or D-amino acids can further reduce immunogenicity and enhance degradation resistance [2].

Thanks to these benefits, cyclic peptide-based therapeutics have garnered significant interest over the past two decades. Currently, over 40 such drugs are used for various applications, including as antibacterial (*daptomycin*), antifungal (*caspofungin*), and immunosuppressant (*cyclosporine A*) agents. Notably, more than 80% of these drugs originate from natural sources or are their derivatives, and very few contain non-canonical or D-amino acids [6]. To expand the diversity of such drugs, it is of great benefit to design cyclic peptides *de novo*.

The current state-of-the-art protein computational design approaches are: (i) machine learning (ML) driven methods, such as those that employ deep neural networks like AlphaFold and RoseTTAFold [7–9]; (ii) physics-based methods, such as those implemented in the Rosetta, Osprey, and Schrödinger software packages [10–12]. Recently, ML methods like AlphaFold have been employed to design sequences for fixed cyclic backbones sampled through physics-based approaches, as well as to propose novel cyclic backbones and sequences, both by using gradient descent methods with a suitably adjusted loss function. For fixed backbone designs, the loss function measured the disparity between the desired backbone and the predicted structure; the goal was to find the sequence minimizing this difference. When creating new cyclic peptide backbones, the loss function was chosen to maximize both prediction confidence and residue interactions. The authors reported that AlphaFold has designed cyclic peptides of 7-13 residues without additional cross-links, and one design per size have been validated experimentally; however, this work has not yet undergone peer review (in preprint [7]).

A major limitation of this ML-based approach is that AlphaFold's neural network training dataset consisted only of canonical L-amino acids. This renders it incapable of designing cyclic peptides composed of mixed L- and D-amino acids, or of integrating non-canonical amino acids. Given the paucity of training data for heterochiral or non-canonical peptides, no existing pure ML approach is likely to prove effective at designing mixed-chirality or non-canonical cyclic peptides *de novo* [13].

Physics-based approaches, which could be better generalized to new chemical building-blocks and new backbone geometries never seen before, are therefore more attractive for heterochiral and non-canonical design. In Rosetta, cyclic backbone conformations are typically sampled using a generalized kinematic closure algorithm (GenKIC) [14, 15]. For an n-residue backbone, torsion angles of $n-3$ residues are sampled randomly, biased by the conformational preferences of the residues, and torsion angles of the remaining 3 residues are solved algebraically to ensure cyclicity. The sampled cyclic backbones are relaxed using energy models that take into account atom pair interactions and torsion angle preferences [16]. Next, Rosetta employs a Monte Carlo simulated annealing algorithm to design an optimal sequence for the relaxed backbones, considering both L- and D-amino acids [17, 18].

This approach has been proven effective in achieving comprehensive sampling of macrocycles ranging from 7 to 10 residues [10]. However, as the peptide size grows, the size of the conformational search space expands exponentially, greatly reducing the likelihood of identifying a stable backbone design through random sampling. Past strategies for dealing with larger peptides have included adding disulfide cross-links to further limit the accessible conformational space (used previously for cyclic peptides of 11-26 residues [10, 15]), limiting conformational searches to symmetric conformations (permitting conformational sampling up to 24 residues [19]), or combining chemical crosslinkers with symmetric sampling and secondary structure biases (which allowed design of a cyclic 60-mer [20]). In each case, sampling of larger structures has been achieved by reducing the generality of the method, and by

imposing more prior expectations of the features present in structures of interest.

In this work, we aimed to overcome the current size limit of general cyclic peptide design, reaching a size of 24 residues with mixed chirality. No additional cross-links or expectations about symmetry or secondary structure needed to be imposed. Our design pipeline, CyclicChamp, consists of (i) sampling "good" cyclic backbones, (ii) optimizing amino acid sequences to align with the backbones, and (iii) validating folded structures of the sequences. Specifically, we focused on steps (i) and (iii), and we followed Rosetta's physics-based approach for our design pursuits. Note that the two approaches by Rosetta and AlphaFold are not mutually exclusive [21], and CyclicChamp could be used for backbone conformational sampling in conjunction with any method, physics- or ML-based, that can carry out sequence design.

We designed cyclic peptides of 7, 15, 20, and 24 residues. For 7 residues, our CyclicChamp yielded high-quality stable designs similar to those designed and experimentally validated by Hosseinzadeh *et al.* [10]. For 15-24 residues, we validated designs through the use of microsecond molecular dynamics (MD) simulations, an algorithmically independent validation approach. These MD simulations generated stable trajectories that indicated promising designs, which were further tested by replica exchange molecular dynamics (REMD) simulations. Two 15-residue, one 20-residue, and one 24-residue designs exhibited thermodynamic stability in REMD simulations, marking them as candidates for future experimental exploration. Importantly, our method is a general one that has permitted us to design large peptide folds that are not dependent on disulfide bonds or other chemical cross-links, predefined symmetry, or human-imposed secondary structure.

Materials and methods

The overall CyclicChamp design workflow is as follows (Fig 1a):

1. We generate a pool of n-residue polyglycine chains, whose initial torsion angles are sampled from a permissive, flattened glycine Ramachandran distribution (which permits conformations accessible to both L- and D-amino acids to be sampled).

2. For each chain, a variant of simulated annealing is performed to search for low-energy configurations that satisfy the cyclic and hydrogen bond (H-bond) constraints.

3. We select representative low-energy configurations for relaxation and sequence design using Rosetta's FastRelax [22] and FastDesign [15, 23], respectively.

4. Low-energy sequences are tested for stability by generating energy landscapes.

In the following subsections, we first derive an error function E_{cyc} for the cyclic backbone constraint (Cyclic error function), and provide relevant energy functions for later backbone sampling (Backbone energy functions). Second, we describe our layered simulated annealing algorithm for low-energy cyclic backbone sampling (Efficient backbone sampling). Finally, we show the two stability analysis methods developed for cyclic peptides of different sizes (Stability analysis for small macrocycles and Stability analysis for large macrocycles).

Cyclic error function

When tackling the backbone closure problem, we consider only the N, C^α, and C' atoms in each residue. A peptide backbone structure is determined by its bond lengths, bond angles, and torsion angles. Bond lengths are the most rigid [24], and are often treated as fixed values as shown in Fig 1b: bond N-C^α with length $d_N = 1.458$ Å, bond C^α-C' with length $d_{C^\alpha} = 1.524$ Å, and bond C'-N with length $d_{C'} = 1.329$ Å. Bond angles and the ω torsion angle can vary to a limited extent ($\pm 5\%$) [24]. The ideal values are $\theta_N = 121.7°$ for the bond angle at atom N, $\theta_{C^\alpha} = 111.2°$ at atom C^α, $\theta_{C'} = 116.2°$ at atom C', and $\omega = 180°$. To simplify computations, we set all these bond lengths and angles to their ideal values.

Fig 1. CyclicChamp workflow and peptide backbone annotations. (A) There are four steps in CyclicChamp for designing stable cyclic peptides. (B) Ideal backbone bond lengths and bond angles are assumed in CyclicChamp. For backbone closure, we consider local coordinate systems at each atom i. There are three steps transforming from coordinate system i to $i+1$.

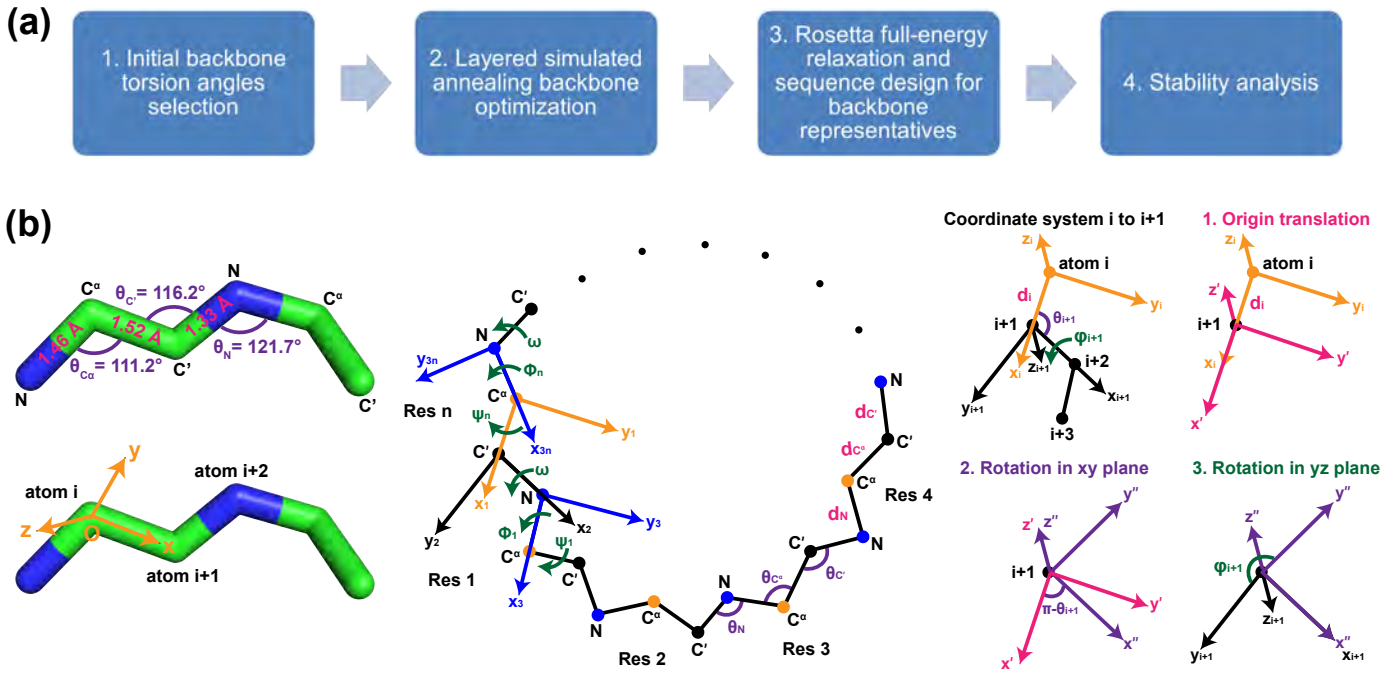

The variables are then torsion angles ϕ and ψ. To close the backbone, Go et al. showed that the torsion angles need to satisfy six independent relations [25]. Following this approach, for an n-residue peptide, we define $3n$ local coordinate systems corresponding to the $3n$ backbone atoms. In coordinate system i, the origin is set to the position of atom i; the x-axis extends towards atom $i+1$; the y-axis is perpendicular to the x-axis and the first quadrant of the xy plane contains atom $i+2$; the z-axis is orthogonal to x- and y-axis using the right-hand rule (Fig 1b). The bond length from atom i to $i+1$ is denoted as d_i. The bond angle formed by atoms $i-1$, i, and $i+1$ is denoted as θ_i. The torsion angle between atoms $i-1$, i, $i+1$, and $i+2$ is denoted as φ_i.

To go from coordinate system i to $i+1$, we need an origin translation $\boldsymbol{p_i}$ of length d_i from atom i to $i+1$, a counterclockwise xy plane rotation by $\pi - \theta_{i+1}$, and a counterclockwise yz plane rotation by torsion angle φ_{i+1} (Fig 1b). A point with coordinate $\boldsymbol{r_{i+1}}$ in system $i+1$ has coordinate $\boldsymbol{r_i}$ in system i, following

$$\boldsymbol{r_i} = \boldsymbol{T}_{\theta_{i+1}} \boldsymbol{R}_{\varphi_{i+1}} \boldsymbol{r_{i+1}} + \boldsymbol{p_i}, \tag{1}$$

where

$$\boldsymbol{T}_{\theta_{i+1}} = \begin{bmatrix} cos(\pi-\theta_{i+1}) & -sin(\pi-\theta_{i+1}) & 0 \\ sin(\pi-\theta_{i+1}) & cos(\pi-\theta_{i+1}) & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

$$\boldsymbol{R}_{\varphi_{i+1}} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & cos(\varphi_{i+1}) & -sin(\varphi_{i+1}) \\ 0 & sin(\varphi_{i+1}) & cos(\varphi_{i+1}) \end{bmatrix}, \boldsymbol{p_i} = \begin{bmatrix} d_i \\ 0 \\ 0 \end{bmatrix}.$$

Because atom 1 is the same as atom $3n+1$ in an n-residue cyclic backbone, backbone closure is then equivalent to having the coordinate systems corresponding to atom 1 and atom $3n+1$ be identical, *i.e.*, same origins and $\boldsymbol{x}$, $\boldsymbol{y}$ directional vectors (rotations preserve dot products, so $\boldsymbol{z}$ is also the same). For

later matrix expression, we choose system 1 to be at atom C^α of residue n (Fig 1b), with origin being $\mathbf{0}$, $\boldsymbol{x}$ directional vector $\boldsymbol{e_1} = [1, 0, 0]^T$, and $\boldsymbol{y}$ directional vector $\boldsymbol{e_2} = [0, 1, 0]^T$.

The origin of system $3n + 1$ has coordinate $\boldsymbol{r_{3n+1}} = \mathbf{0}$ in system $3n + 1$ and $\boldsymbol{r_1}$ in system 1. To satisfy backbone closure, we require

$$\boldsymbol{r_1} = \boldsymbol{M_1 M_2 \cdots M_{n-2} M_{n-1} q} + \boldsymbol{M_1 M_2 \cdots M_{n-2} q} + \cdots + \boldsymbol{M_1 M_2 q} + \boldsymbol{M_1 q} + \boldsymbol{q} = \mathbf{0}, \tag{2}$$

where $\boldsymbol{M_i} = \boldsymbol{T_{\theta_{C'}} R_\omega T_{\theta_N} R_{\phi_i} T_{\theta_{C^\alpha}} R_{\psi_i}}$, and

$$\boldsymbol{q} = \boldsymbol{T_{\theta_{C'}} R_\omega T_{\theta_N}} \begin{bmatrix} d_N \\ 0 \\ 0 \end{bmatrix} + \boldsymbol{T_{\theta_{C'}}} \begin{bmatrix} d_{C'} \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} d_{C^\alpha} \\ 0 \\ 0 \end{bmatrix}.$$

For the $\boldsymbol{x}$ and $\boldsymbol{y}$ directional vectors in system $3n + 1$, we require their vector forms in system 1 equal to $\boldsymbol{e_1}$ and $\boldsymbol{e_2}$, respectively, *i.e.*,

$$\boldsymbol{M_1 M_2 \cdots M_{n-2} M_{n-1} M_n e_1} = \boldsymbol{e_1}, \quad \boldsymbol{M_1 M_2 \cdots M_{n-2} M_{n-1} M_n e_2} = \boldsymbol{e_2}. \tag{3}$$

See the detailed derivation in S1 Appendix.

Combining requirements Eq (2) and Eq (3), and using squared error, we write the cyclic constraint into a single equation,

$$\begin{aligned} E_{cyc} = & ||\boldsymbol{q} + \boldsymbol{M_1 q} + \boldsymbol{M_1 M_2 q} + \boldsymbol{M_1 M_2 M_3 q} + \cdots + \boldsymbol{M_1 M_2 M_3 \cdots M_{n-1} q}||_2^2 \\ & + ||\boldsymbol{M_1 M_2 M_3 \cdots M_{n-1} M_n e_1} - \boldsymbol{e_1}||_2^2 \\ & + ||\boldsymbol{M_1 M_2 M_3 \cdots M_{n-1} M_n e_2} - \boldsymbol{e_2}||_2^2. \end{aligned} \tag{4}$$

We call this the cyclic error function, and finding a cyclic peptide backbone solution is equivalent to finding a zero for Eq (4). Note that all the bond lengths, bond angles, and torsion angle ω are fixed to ideal values, so $\boldsymbol{M_i}(\phi_i, \psi_i)$ are matrices of variables ϕ_i and ψ_i, and vector $\boldsymbol{q}$ can be calculated explicitly for ideal bond angle and bond length values as $\boldsymbol{q} = [3.5620, 1.3322, 0]^T$.

Backbone energy functions

Our backbone sampling algorithm is independent of the energy model, so can work with any energy functions that are fast to evaluate. Here, we choose Rosetta's *Ref2015* energy model [16]. For the backbone atoms, we evaluate the Ramachandran, repulsive (Van der Waals), attractive (London dispersion), electrostatic, solvation, and H-bond energy terms. When analyzing backbone atom pair interactions, we consider only atom pairs separated by at least 4 covalent bonds, as is the case in the *Ref2015* energy function [16]. See detailed descriptions in S2 Appendix and S1 Fig.

Each amino acid has a Ramachandran map that shows the energetically allowed regions in ψ-ϕ space. Glycine has the largest Ramachandran area, proline has the smallest, and the other amino acids have similar areas but different energetically favorable regions. As L- and D-amino acids are mirror images, their accessible Ramachandran regions are mirror-symmetrical. For backbone sampling, we use the permissive, symmetrized glycine Ramachandran map to allow all possible amino acids for later sequence design (S2 Fig, regions within the blue boundary). This map is based on the statistical distribution of glycine conformations observed in Protein Data Bank structures but made symmetric as previously described [10, 15].

Efficient backbone sampling

Initial torsion angles selection

We partition the glycine Ramachandran map into six torsion bins, marking bin centers (S2 Fig). For each residue, its initial angles ϕ and ψ are chosen randomly from one of the six centers. For an n residue

peptide, there are 6^n initial point combinations (or initial configurations). Considering equivalence classes induced by cyclic permutations (*e.g.*, "1232456" is equivalent to "2324561"), the number of unique combinations is reduced to $\frac{1}{n}\sum_{i=1}^{n} 6^{gcd(i,n)}$, where $gcd(i,n)$ is the greatest common divisor of i and n; this is approximately $6^n/n$ for large n [19, 26]. For 7 residues, this yields 39,996 initial point combinations. As the macrocycle size increases, the number of combinations increases exponentially. Hence, for large macrocycles with 15, 20, or 24 residues, we randomly select 100,000 initial point combinations.

Layered simulated annealing

With the initial angles assigned, we search the Ramachandran space to find low-energy configurations satisfying the cyclic and H-bond constraints. To save computational cost, we devised a simulated annealing (SA) variant with multiple layers of acceptance criteria. At time step t, a random move within a disk of radius k_t in the Ramachandran space is generated for each residue (S2 Fig). If a move enters a prohibited high-energy region of Ramachandran space (white area in S2 Fig), it is rejected for that residue. The subsequent new configuration needs to pass four layers of energy tests to be accepted: sequentially, Ramachandran energy, repulsive energy, cyclic error, and H-bond energy. For 7-residue peptides, a final energy test is added to calculate miscellaneous (attractive, electrostatics, and solvation) energies.

At each test layer l, we use the Metropolis criterion, *i.e.*, the new configuration passes if it has an energy $E_{new,l}$ lower than the current energy E_l, or below a threshold $E_{thr,l}$. If neither holds, the new configuration has a probability of $e^{(E_l - E_{new,l})/T_{t,l}}$ to pass the test, where $T_{t,l}$ is the temperature at time step t and test l. Once a new configuration passes all tests, it is accepted and becomes the current configuration. Note that this approach is not intended to produce thermodynamic distributions of states as pure Metropolis-Hastings Monte Carlo trajectories do; instead, the goal is to rapidly discover low-energy states with the least computational expense needed, by evaluating cheaper energy terms first to reject moves.

Configurations that have low repulsive energy, low cyclic error, and strong H-bonds are recorded as good backbone candidates. Details for choosing the energy thresholds, good backbone criteria, and other simulated annealing parameters are provided in S3 Appendix. In particular, within the thousands of possible simulated annealing parameter combinations, we use combinatorial design [27] to select and test 51 combinations for 7- and 15-residue tests, and 400 combinations for 20- and 24-residue tests (details in S4 Appendix).

Once we find the optimal simulated annealing parameter combination (S1 Table), for each initial configuration, we run the layered simulated annealing algorithm. If a single run of simulated annealing does not produce any good backbone candidates, we repeat again, for a maximum of three repeats. Two example runs that successfully produced good 15-residue backbone candidates are uploaded to GitHub for illustration.

Backbone clustering and sequence design

There can be vast numbers of backbone candidates, and many are similar to each other. To select lowest-energy representatives, we clustered the candidates based on the torsion bins in a manner similar to that previously described [10]. Briefly, we assigned a torsion bin number to each residue in a backbone candidate (S2 Fig), to produce a torsion bin string. For example, a string "1351246" means that the first residue in the backbone falls in torsion bin 1 of the glycine Ramachandran space, the second residue in torsion bin 3, the third residue in torsion bin 5, and so on.

We considered torsion bin strings as equivalent if they can be cyclically permuted, such as "1351246" and "3512461". To uniquely identify the equivalence class, we looked for the cyclic permutation that moves the smallest bin value to position 1. If multiple residues share the same smallest value, we chose the permutation that gives smaller value at position 2, and so on. In this way, we clustered the backbone candidates into equivalence classes of torsion bin strings. Within each class, we selected the candidate with the lowest energy as the representative.

We computed energies based on Rosetta $Ref2015$'s weights [16],

$$E_{total} = 0.45 * E_{rama} + E_{rep} + E_{hbond} + E_{other}.$$

These backbone representatives were then sent for full-energy relaxation using Rosetta's FastRelax, following past protocols [4, 10, 19] (scripts in S5 Appendix and uploaded to GitHub). Relaxed backbones with low energies were selected for further sequence design using Rosetta's FastDesign, permitting the 20 canonical amino acids (except cysteine and glycine) with both their L and D forms [10] (S6 Appendix). For 20- and 24-residue designs, to avoid instability caused by buried unsatisfied polar atoms, additional restrictions of amino acid types were applied [19] (S7 Appendix).

Stability analysis for small macrocycles

Designed sequences having low energies underwent final stability analysis. To assess the stability of a designed sequence, we sampled alternative conformations for this sequence. The energies of these alternative conformations, together with their root-mean-square-deviations (RMSDs) from the designed structure, form the energy landscape. To calculate RMSD, we used the Kabsch algorithm [28] to align backbone heavy atoms (N, C^{α}, C', and O) of an alternative conformation with those of the designed structure.

If the lowest-energy conformations all have small RMSDs from the designed structure, then the sequence has a high chance to fold into the designed structure, and we consider the design stable. In order to provide a quantitative measure of stability, we employ the P_{Near} value introduced in 2016 [15]:

$$P_{Near} = \frac{\sum_{i=1}^{N} e^{-\frac{RMSD_i^2}{\lambda^2}} e^{-\frac{E_i}{k_B T}}}{\sum_{j=1}^{N} e^{-\frac{E_j}{k_B T}}}, \tag{5}$$

where $k_B T = 0.62$ kcal/mol (equivalent to 37 °C), $\lambda = 0.5$ for small macrocycles (7 residues), $\lambda = 1.5$ for medium macrocycles (15 residues), and $\lambda = 2$ for large macrocycles (20 and 24 residues). It has been experimentally shown that a $P_{Near} > 0.9$ is indicative of stability, and correlates well with experimental success in binder design [4, 5, 10].

For macrocycles having 7 residues, our backbone simulated annealing algorithm considered all combinations of initial angle torsion bins, so backbone sampling was comprehensive. Hence, we can expect that all low-energy conformations exist within the sampled backbones. By threading designed sequences on these backbones and using Ramachandran map at each residue as a series of filters to weed out incompatible conformations, we were able to approximate the energy landscape containing all alternative low-energy conformations. Given that the backbones are cyclic, it was necsesary to examine all n permutations of the residues for each backbone conformation. Compatible backbones were then subjected to FastRelax, following the same protocol (S5 Appendix), but this time with the designed sequence instead of poly-glycine. We refer to this energy landscape sampling process as *Ramachandran-stability filtering.*

Stability analysis for large macrocycles

For macrocycles of 15-24 residues, the Ramachandran-stability filtering method failed to find a sufficient number of alternative backbone conformations due to the exponentially larger search space (see S8 Appendix). Simply adapting the layered simulated annealing algorithm for designed sequences was not enough. We aimed to explore energy landscapes filled with local minima, spanning both low and high RMSD regions (0-6 Å). Simulated annealing might sometimes jump across certain minima without sufficient exploration, or conversely, become trapped in some minima without investigation of others. To overcome this problem, we employed a genetic algorithm. This approach allowed us to broadly explore the landscape in the initial stages, and through successive generations, to probe and settle into the low-energy minima, thus achieving a balanced exploration of both global and local features of the energy landscapes.

The initial population of the genetic algorithm comprised alternative structures of the designed sequence, whose backbones were sampled by two separate layered simulated annealing trajectories, one targeting low energy and, when a designed backbone was available, the other targeting low RMSD. The low energy simulated annealing protocol was as described in S3 Appendix, while the low RMSD simulated annealing had only the cyclic error test and a RMSD test (see S8 Appendix). In the RMSD test, we used the Kabsch algorithm [28] to measure the backbone-heavy-atom RMSD between the new configuration and the designed structure, and used the Metropolis criterion to accept new configurations.

The sampled backbones were then subject to FastRelax, with the designed sequence specified so that the corresponding side-chains were added and optimized by FastRelax (S5 Appendix). We sorted the relaxed structures in ascending order of their energies and initiated energy-based clustering (see S8 Appendix). The $2 * N_{GA}$ lowest-energy cluster centers formed the initial genetic algorithm population.

In the genetic algorithm, each generation underwent crossover, mutation, and selection. Crossover involved checking whether a pair of parents could exchange residues within a designated region. Mutation involved random perturbation of torsion angles within a given region. See details in S8 Appendix and S4 Fig. After collecting the crossover and mutation children, we ran FastRelax to add side-chains and obtain full energies. Due to distortions of bond angles and lengths caused by crossover and mutation at breakpoints, we used Cartesian relaxation to restore near-ideal bond geometry (scripts provided in S9 Appendix). Then, we clustered the relaxed structures and selected the lowest-energy N_{GA} cluster centers to form the next generation.

During each generation, of all the cluster centers, we recorded those having energies < 0 for the eventual energy landscape. We conducted 50 such generations with $N_{GA} = 500$ for 15 residue macrocycles, $N_{GA} = 750 - 5 \cdot i$ for generation i of 20 residues, and $N_{GA} = 1000 - 10 \cdot i$ for generation i of 24 residues. The reduced population size in later generations was for efficiency: a broad exploration of the energy landscape is beneficial for early stages, while in later stages the genetic algorithm focuses on exploring the minima, which doesn't require a large population. We refer to this energy landscape sampling algorithm as *ClusterGen*.

Molecular dynamics simulations for top designs

As a means of validating designs independent of Rosetta or the conformational sampling methods developed here, we carried out molecular dynamics (MD) simulations. For 15-24 residue designs that had high P_{Near} scores, we performed 1-μs MD simulations to validate kinetic stability. We used *OpenMM v8.1.0.beta* [29] toolkit with the *amber14/protein.ff14SB* force field for peptide and the *amber14/tip3p* force field for water. The water box had periodic boundary conditions and 1 nm padding distances from the peptide. The ionic strength was set to be 0.15 molar with Na+ and Cl-. Full details are in S10 Appendix.

We also ran replica exchange molecular dynamics (REMD) simulations [30] for top designs that showed stable MD trajectories, using *OpenMMTools v0.23.1* [31] (details in S10 Appendix). After an initial 100 ns simulation, we plotted the radius of gyration (Rg) distributions for uncorrelated configurations extracted by *OpenMMTools MultiStateSamplerAnalyzer* [31] from two different time intervals 50-70 ns and 80-100 ns. If the distributions do not overlap well, we extended the REMD for an additional 50 ns.

To compute the average RMSD of a simulation, we extracted uncorrelated configurations sampled at 300 K, and calculated their C^{α}-atom RMSDs from the initial design using *MDTraj v1.9.8* [32]. Free energy surfaces (FES) were derived using RMSD and Rg as collective variables. We gathered RMSD and Rg data from 300 K uncorrelated configurations, and computed the probability densities $P(RMSD, Rg)$ using the *histogram2d* function in the Python package *numpy* [33], with 50 bins along each dimension. The free energy was calculated as $-RT\ln[P(RMSD, Rg)]$, where R is the gas constant and T is the temperature 300 K. Note that when calculating the probability densities, we used the standard unit nanometer for RMSD and Rg, yet for visualization consistency, we plot the FES in Å.

We used the Flatiron Rusty cluster GPU nodes for the MD simulations. Each node was equipped with four NVIDIA 40 GB A100 Tensor Core GPUs (Ampere), 1024 GB system memory, and 64 CPU

cores. We ran each MD or REMD simulation on one A100 GPU. Typical MD runs took about 40 hours to complete, while typical 100 ns REMD runs took about 5-7 days to complete.

Results

Our cyclic peptide design pipeline CyclicChamp consists of several key steps: initial backbone torsion angle selection, backbone sampling through layered simulated annealing, backbone clustering using torsion bin strings, backbone relaxation with FastRelax, sequence design via FastDesign, and stability analysis by generating energy landscapes (see Fig 1). In Fig 2, we show the CPU-hours required in each step, as benchmarked on New York University's Greene High Performance Computing Cluster. A more detailed ClusterGen computation time breakdown is presented in S5 Fig, and the number of candidates generated in each step is listed in S6 Fig.

Additionally, we compared the design results for 7 and 15 residues with Rosetta by running the

Fig 2. CyclicChamp computation time and design comparisons with Rosetta. (A) The computation time required by CyclicChamp backbone sampling and stability validation (ClusterGen) exhibits linear-like growth with increasing backbone size. FastDesign was faster for 20 and 24 residues than for 15 because there were fewer backbones on which we did sequence design. (B) Total design time divided by the number of stable designs validated by the filtering method for 7 residues, ClusterGen for 15 residues, and reshaped ClusterGen for 20 and 24 residues. (C) When allocating equivalent computation time for backbone sampling, CyclicChamp generated 5 to 7 times as many cyclic backbones with sufficient H-bonds as Rosetta's *simple_cycpep_predict*, which led to double or triple number of stable designs as Rosetta's after stability validation. (D) Rosetta was not able to scale up to asymmetric designs of size 20 and 24, so we show the design statistics only for CyclicChamp.

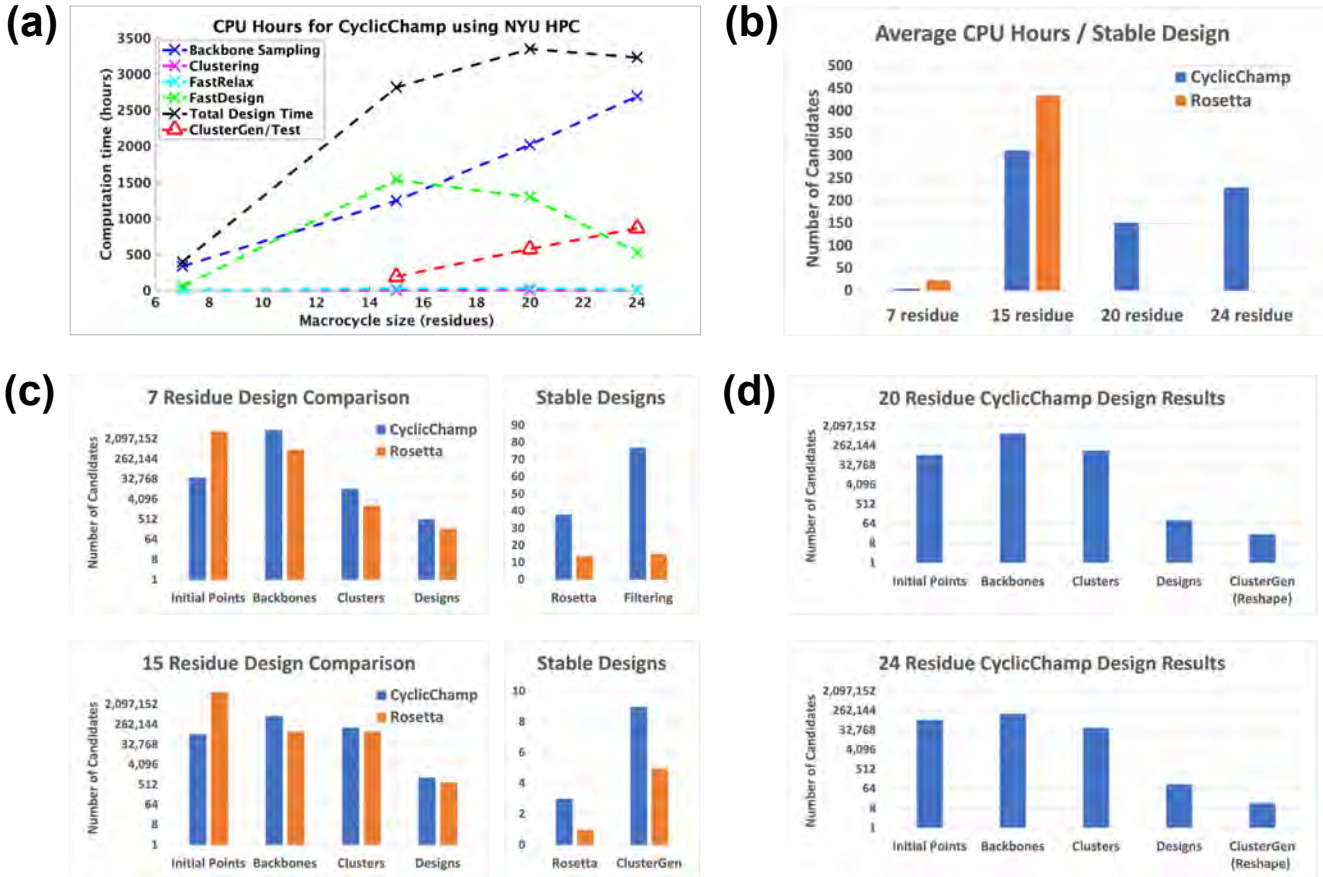

design pipeline protocol described in [10], and this highlighted the greater efficiency of our new method. We did not run Rosetta for 20 and 24 residue designs due to its low probability of hitting a stable design by random sampling, as suggested by the energy landscape generation in the stability analysis below (Fig 6b, Fig 7b). For all designed sequences, no low-energy conformation could be found by *simple_cycpep_predict* with 100,000 random attempts.

The computation time for backbone sampling exhibited linear-like growth as the size of the backbone increases (Fig 2a). The time growth from 7 to 15 residues primarily resulted from a rise in the number of initial backbone configurations, escalating from 39,996 to 100,000. Beyond 15 residues, we fixed the number of initial configurations, so that the time growth was largely due to the $\mathcal{O}(N^2)$ complexity involved in calculating atom pairwise energies (see S2 Appendix). However, the sampling of 24-residue backbones yeilded only about one-fourth the number of backbone clusters compared to those from 15 and 20 residues, indicating a substantial increase in the difficulty of finding good backbone candidates (S6 Fig).

We conducted P_{Near} stability analyses on low-energy designs using our Ramachandran-stability filtering method for 7 residues and our Clustering genetic algorithm (ClusterGen) for 15-24 residues. To obtain more stable 20- and 24-residue cyclic peptide structures, we also looked at alternative low-energy structures sampled by ClusterGen and reshaped the energy landscapes accordingly (see details in the section titled Large macrocycle 20 residue designs). The average design time required for obtaining a stable design with $P_{Near} > 0.9$ is plotted in Fig 2b. Notably, our CyclicChamp operated more efficiently, especially for 7-residue designs, requiring only one-fourth the time of Rosetta's to find a stable design. The reduction in CyclicChamp design time when going from 15 residues to 20 residues was attributed to a less stringent stability analysis, *i.e.*, the inclusion of alternative low-energy structures sampled by ClusterGen.

Our filtering method took about 1.3 CPU-hours per stability test. The computation time of ClusterGen grew linearly as the macrocycle size increased, primarily due to the progressively larger populations adopted (see Stability analysis for large macrocycles). Rosetta's *simple_cycpep_predict* was also used as a stability test, but it failed to adequately explore the exponentially larger conformational space as the size went up to 20 and 24.

Looking at the specific steps in the design pipeline, we found that CyclicChamp had greater efficiency in backbone sampling (Fig 2c). Within the allocated 336 CPU-hours for 7-residue backbone sampling and 1248 hours for 15 residues, CyclicChamp found five to seven times as many cyclic backbones with sufficient H-bonds. After clustering, the lowest-energy half of the cluster centers were advanced to sequence design. CyclicChamp achieved approximately twice the number of designs with energies below the thresholds of -8 kcal/mol for 7 residues and -30 kcal/mol for 15 residues compared to Rosetta. After stability validation, CyclicChamp managed to produce two to three times the number of stable designs ($P_{Near} > 0.9$) as compared to Rosetta, illustrating its superior capability in finding high-quality cyclic peptide backbone conformations that are more likely to result in stable designs.

In the following sections, we provide stability analysis results for our 7-24 residue macrocycle designs, as well as the molecular dynamics simulation validations.

Small macrocycle 7 residue designs

For the top 513 designs with energies below -8 kcal/mol, we compared P_{Near} values produced using Rosetta's *simple_cycpep_predict* application and with our Ramachandran-stability filtering method (Fig 3a). The Pearson correlation coefficient stood at $r = 0.822$, and 94% of the values exhibited deviations less than 0.4. There were 38 designs having $P_{Near} > 0.9$ by the Rosetta's method, 77 designs having $P_{Near} > 0.9$ by the Ramachandran-stability filtering method, and of these 32 had $P_{Near} > 0.9$ by both methods.

We also analyzed the relationship between P_{Near} values and energies (Fig 3b). We divided the 513 designs into four evenly spaced energy bins from -16 to -8 kcal/mol. Within each energy bin, we calculated the relative frequency of P_{Near} values in three ranges: low ($P_{Near} < 0.7$), medium ($0.7 \leq P_{Near} < 0.9$), and high ($0.9 \leq P_{Near}$). For both methods, we found that the second-lowest energy bin possessed the largest relative probability of high P_{Near} values, suggesting that the the

Fig 3. Stability analysis of 7-residue designs. (A) Correlation plot between P_{Near} values calculated by Rosetta *simple_cycpep_predict* and by our filtering method. (B) P_{Near} value distributions within different energy bins (kcal/mol). Design counts are labeled on top of the bars. (C) Energy landscape comparisons for three cases that have noticeable differences in P_{Near} values calculated by the two methods. Left, more extensive sampling of wells further from the designed state by the filtering method results in a lower computed P_{Near} value. Middle: more extensive sampling close to the designed state by the filtering method identifies a deeper minimum, raising the P_{Near} value. Right: more extensive sampling by the filtering method allows exploration in a low-RMSD region missed entirely by Rosetta, identifying energy wells and raising the P_{Near} value. (D) Top 7-residue designs ($P_{Near} > 0.9$) demonstrating smaller backbone root-mean-square radii with H-bond intersections. Representative designs having 0, 1, 2, 3, 5, and 6 H-bond intersections are drawn, along with their H-bond networks where L- and D-amino acids are specified and arrows point from amide proton to carbonyl oxygen.

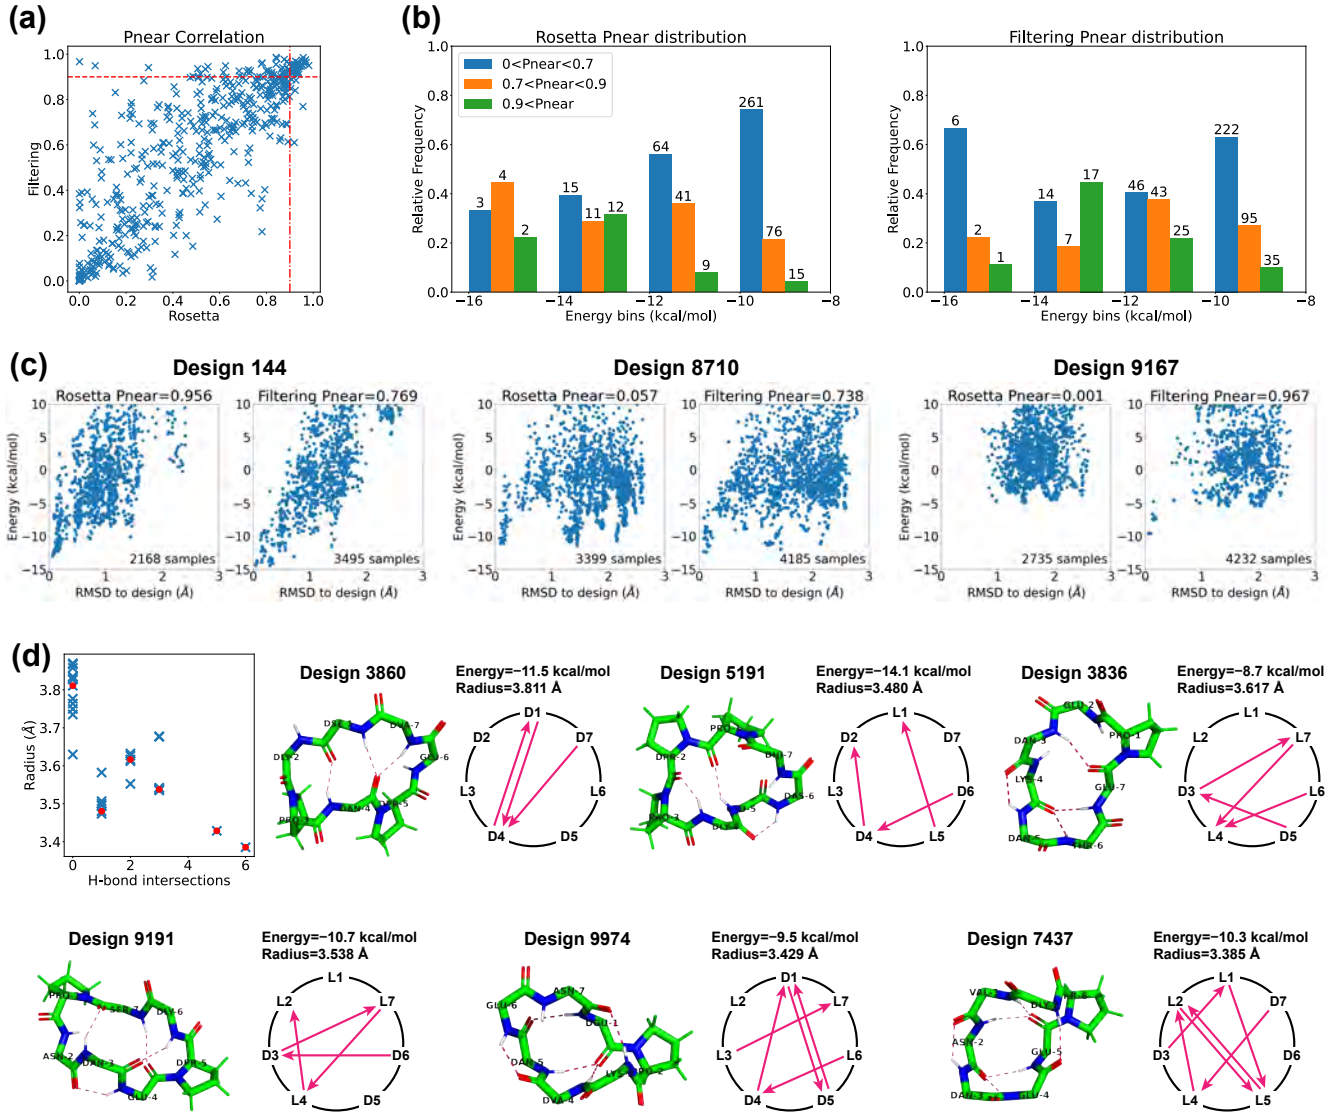

conformational energy of the designed state should not necessarily be as low as possible in order to achieve a design with a large energy gap between the designed state and all alternative states. This may be attributed to favourable but non-specific interactions stabilizing both the designed state and alternative states for the designs in the lowest-energy bin. We noticed that one high P_{Near} design in the

Fig 4. Comparison with Rosetta experimentally validated 7-residue designs [10]. (A,B) From our 513 designs, we find designs (colored in orange) for which the backbones align best with the Rosetta designs (light blue). The residues having different side-chains are marked in red. (C) Design 475 has an alternate low-energy structure (purple), leading to a low P_{Near} value.

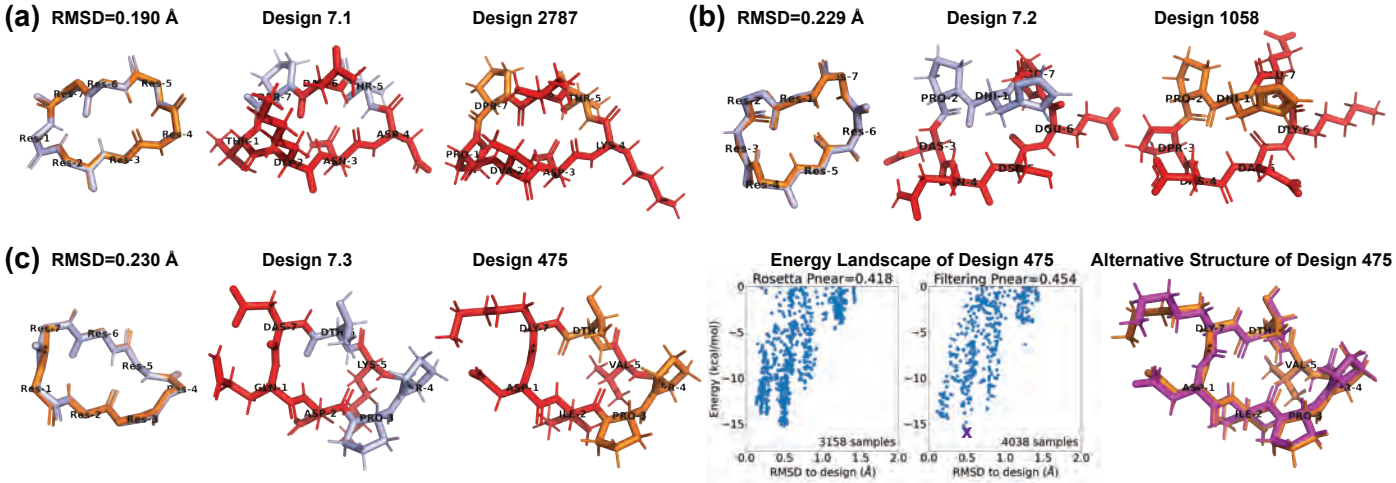

first energy bin dropped from 0.956 to 0.769 when switching from the Rosetta's method to the filtering method (Fig 3c, design 144). It appears that in this case the filtering method was able to sample the low RMSD region more comprehensively than Rosetta could, identifying alternative low-energy structures around 0.5 Å and 1.5 Å RMSD from the design whose populations reduce the fractional occupancy very close to the designed state, lowering P_{Near}.

Additionally, we examined designs in which the P_{Near} values of the two methods differed by more than 0.4. In most cases, the two methods yielded energy landscapes with similar shapes but different distributions (Fig 3c, design 8710). Nevertheless, there were instances in which the Ramachandran-stability filtering method alone proved capable of sampling the low RMSD region (design 9167). While we have never seen a case in which Rosetta samples low RMSD regions that the Ramachandran-stability filtering method misses, there is no way to check this exhaustively. Operationally, we suggest using both methods.

Overall, the Ramachandran-stability filtering method holds a distinct advantage over Rosetta's random sampling approach in scrutinizing the low RMSD region and unearthing alternative low-energy structures. This is likely due to its strategy of selecting sample backbones from the candidate pool created by layered simulated annealing, characterized by their low repulsive energies and robust H-bonds. By randomly sampling the torsion angles from their Ramachandran spaces except for three residues, which are solved algebraically to ensure cyclicity through generalized kinematic closure (GenKIC) [14] Rosetta's *simple_cycpep_predict* can sometimes fail to find the low RMSD region, especially as the dimensionality increases.

For the 32 designs that have $P_{Near} > 0.9$ by both methods, we observed an interesting correlation between the structural compactness and the H-bond patterns. We measured the structural compactness by calculating the root-mean-square distance/radius of the backbone atoms from the backbone center of mass. When drawing the H-bond networks, we noticed that some designs have intersecting backbone H-bonds, and such designs tend to have smaller backbone radii (Fig 3d). The H-bond intersection counts varied from 0 to 6, and we show one representative design for each count. Design 7437 had the smallest radius and adopted a semi-ball shape due to its six intersecting H-bonds. Note that the compactness was not simply a product of more H-bonds. For example, despite all having three H-bonds, design 5191 with one H-bond intersection presented a radius of 3.480 Å, while design 3860 with no intersection exhibited a radius of 3.811 Å. All the 32 design structures and their associated information are uploaded to GitHub.

Finally, we compared our designs with the previously-published comprehensive design results from Rosetta [10] to see whether similar structures have been found. For 7 residues, that study reported 12 Rosetta-produced designs with $P_{Near} > 0.9$, of which three were experimentally validated (Design 7.1-3 in Fig 4). We aligned the backbones of our 513 designs with these 12 Rosetta designs. For each pair, we tried all seven cyclic permutations to find the best alignment. Our closest backbones have 0.114-0.639 Å RMSDs from the Rosetta design backbones, suggesting that our layered simulated annealing algorithm can find similar backbone conformations that can be stabilized with a suitable choice of sequence.

Most of these simulated annealing designs had a sequence match of 2-3 residues with their corresponding Rosetta designs. This follows the observation in the 2017 Rosetta study that usually fewer than three residues (often prolines) are critical for maintaining the fold [10]. Additionally, all sequences retained the same pattern of chirality (L or D-amino acids) as the Rosetta designs, with one exception that exhibited a low $P_{Near} < 0.3$ according to both Rosetta and the Ramachandran-stability filtering method. This was in stark contrast to most other designs with $P_{Near} > 0.7$, aligning with the notion that altering chirality is more disruptive than replacing an amino acid with another of the same chirality [10].

For the three experimentally validated Rosetta designs (7.1-7.3), we found closely matched simulated annealing designs with ∼0.2 Å backbone RMSDs (designs 2787, 1058, and 475 in Fig 4). At least one proline was preserved in each pair. Designs 2787 and 1058 maintained stable folds with $P_{Near} > 0.9$, while design 475 had only $P_{Near} \sim 0.4$. We plotted the energy landscapes for design 475, and found an alternative low-energy structure in the 0.5 Å RMSD region (Fig 4c). The main deviations between design 475 and its alternative structure were at residues 1 and 7, consistent with the observed turn flip around residue 7 in the NMR structural ensemble of Rosetta design 7.3 [10].

Medium macrocycle 15 residue designs

Past Rosetta studies have used disulfide cross-links to design cyclic peptides of 11-14 residues [10] and structural symmetry to design larger sizes [19, 20]. Our methods were able to design general stable, computationally-validated 15-residue cyclic peptides. For P_{Near} stability analysis, we used both Rosetta's *simple_cycpep_predict* and our ClusterGen. In the energy landscapes, the use of Cartesian coordinate relaxation leads to an energy reduction of approximately 15 kcal/mol compared to designs using torsion angle relaxation (Fig 5b,d). Because of the high computational cost of validation (100,000 samples for *simple_cycpep_predict*, and 50 generations of a population of 500 for ClusterGen), we validated only the top 75 designs having the lowest energies.

A correlation plot in Fig 5a compares the P_{Near} values computed by ClusterGen and Rosetta. The P_{Near} values from ClusterGen were evenly distributed across all ranges, and there were nine designs with $P_{Near} > 0.9$. Meanwhile, 64% of Rosetta's P_{Near} values fell below 0.1, and only three designs achieved $P_{Near} > 0.9$. These three designs also exhibited high P_{Near} values according to ClusterGen, and their energy landscapes' low RMSD regions were better explored by ClusterGen, as seen in Fig 5b. In instances where significant differences existed between the two methods' P_{Near} values, Rosetta tended to underestimate P_{Near} due to failure to extensively sample low RMSD regions (Fig 5c). In contrast, ClusterGen consistently sampled the low RMSD region, and distinguished peptides by generating energy landscapes of various shapes, such as the dual-minima landscape observed in Design 108020, the broad energy minimum in Design 120510, and the sharp funnel shape in Design 2599. When we plotted ClusterGen's P_{Near} values against the structure compactness measured by backbone radii (Fig 5a), we found that designs with small radii, particularly those under 5.5 Å, tended to exhibit high P_{Near} values.

Interesting backbone structural motifs appeared in the top designs ($P_{Near} > 0.9$ by either Rosetta or ClusterGen). Designs 169032, 2599, and 16897 shared a recurring structural motif in which the backbone bends (Fig 5d). At the bending locations, we saw $i, i+3$ H-bonds, with the CO group in residue i binding to the NH group in residue $i+3$. Such $i, i+3$ H-bonds can cause backbone turns, as observed in the 2017 Rosetta study [10]. In these three designs, the $i, i+3$ turns were paired with extra long-range H-bonds. The CO group of residue $i+1$ or the NH group of residue $i+2$ bound to the opposite backbone side, leading to a simultaneous bend on both sides. Despite the fact that all three designs have intersecting H-bonds, the H-bonds in design 16897 are closely intertwined, while the H-bonds in design

Fig 5. Stability analysis of 15-residue designs. (A) Correlation plot between P_{Near} values calculated by Rosetta *simple_cycpep_predict* and by our ClusterGen. The P_{Near} values of our ClusterGen are also plotted against the backbone root-mean-square radii. (B) Energy landscape comparison for one case in which both methods obtain high P_{Near} values. (C) Energy landscape comparisons for three cases in which the two methods calculate significantly different P_{Near} values. (D) Top designs with bending backbones. The backbone atoms, prolines (in color purple), and hydrophobic amino acids (ALA, ILE, LEU, VAL, MET, PHE in color orange) are shown. The backbone turn segments are enlarged. (E) Top designs with short alpha helices and consecutive $i, i+2/i+3$ H-bonds.

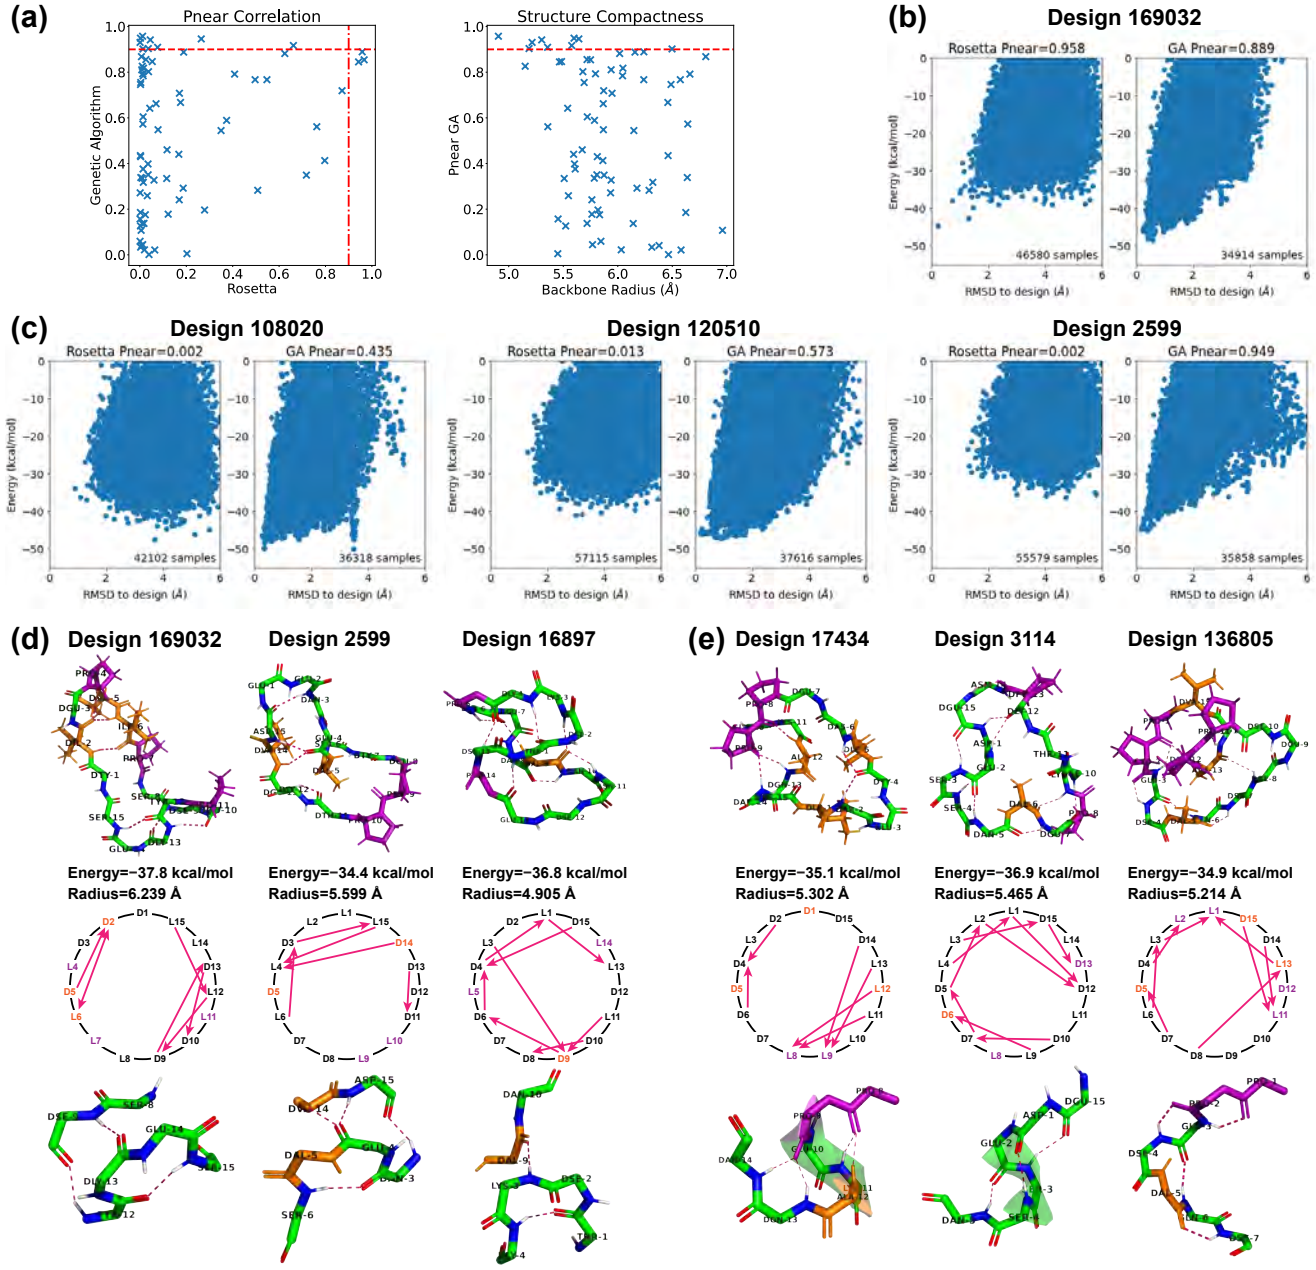

169032 can be distinctly separated into two parts. Consequently, the backbone root-mean-square radius varies from 4.905 to 6.239 Å.

Another three designs contained ordered consecutive H-bonds (Fig 5e). In design 17434, residues 13 and 14 bound to residue 9, and residues 11 and 12 bound to residue 8. These H-bonds held the twisted

backbone tightly to form a short alpha helix. In design 3114, three consecutive $i, i+3/i+4$ turns among residues 15 and 1-5 shaped another short alpha helix. When the $i, i+3$ turn paired with $i, i+2$ H-bonds, there resulted a semi-circular segment as seen in residues 1-7 of design 136805. These H-bond arrangements may prove valuable as building blocks for future design endeavors.

Large macrocycle 20 residue designs

We first validated 22 low-energy designs using both Rosetta and ClusterGen to compare the performance of the two algorithms. The P_{Near} value correlation plot is shown in Fig 6a. Notably, all of Rosetta's P_{Near} values fell below 0.2, while ClusterGen's P_{Near} values spanned evenly from 0.001 to 0.822. This was caused by Rosetta's failure to sample the low RMSD regions, as can be seen in the round-shaped energy landscapes in Fig 6b. We then validated 47 more low-energy designs using only ClusterGen. By plotting the ClusterGen's P_{Near} values against backbone root-mean-square radius, we again saw that designs with high P_{Near} values tended to have small radii (Fig 6a). The highest P_{Near} value obtained was 0.897 with a radius of 6.196 Å.

As our primary goal was to find stable cyclic peptide structures, rather than to design fixed backbone conformations, we also looked at alternative low-energy structures in the energy landscapes of our designs. We reshaped these landscapes by selecting the lowest energy structure sampled as the native state for each designed sequence, and recomputed the backbone RMSD values of all samples from this new native state (Fig 6b). We then calculated the P_{Near} values for these reshaped energy landscapes. Note that ClusterGen is not biased towards the initial designed structure, because in each generation, we choose the lowest energy samples to form the next generation. As long as the original energy landscape is thoroughly explored, the reshaped landscape will reflect the stability of the new native state. This reshaping yielded 22 low-energy structures with $P_{Near} > 0.9$. From these, we selected the 11 with energies below -65 kcal/mol, as well as the top design 35869 that has an original $P_{Near} = 0.897$ before reshaping; these are presented in Fig 6c,d.

Design 35869 had a hydrophobic core (residues 7-11) that formed four backbone H-bonds with the surrounding residues, imparting rigidity. Among the selected low-energy structures, five exhibited local conformational deviations from their initially-designed configurations, so we depict only their low-energy conformations in Fig 6c. Notably, structure 80837 demonstrated a quasi-cyclic (C2) symmetry within its backbone, an intriguing feature considering the sequence's inherent asymmetry. This structure formed three $i, i+3$ backbone turns across residues 11-16 and 3-5, with an H-bond between residues 4 and 14 consolidating a hydrophobic core. Structure 68384 had the most backbone turns, five $i, i+3$ H-bonds highlighted in Fig 6c, resulting in a twisted structure with a small backbone radius. In structure 15036, a 3_{10}-helix was present in residues 13-16, accompanied by two $i, i+3$ turns.

The remaining six low-energy structures displayed considerable conformational shifts from their originally designed configurations, as illustrated in Fig 6d. In structure 45902, a half-turn helix emerged, enhancing the stability of the loose backbone end. Structure 83218 underwent a transformation that introduced a complete helical turn, significantly decreasing the backbone radius from 7.185 Å to 6.125 Å, with a corresponding P_{Near} value increase from 0.035 to 0.924. The other four structures formed dense H-bonding that lead to compactness. Specifically, the consecutive H-bonds between residues 3-7 and 19-20 in structure 1665 pulled the backbone segments closer. In structure 74102, an H-bond network among residues 1 to 3 facilitated a 90-degree turn, effectively altering the backbone from a flat to a more globular shape. Structures 107505 and 25226 each introduced a single long-range H-bond between residues 8 and 14, and residues 7 and 16, respectively, drawing the loose backbone ends together.

These top structures all had nicely packed hydrophobic cores made of 1-5 residues (S7 Fig). The number of proline residues varied from 2 to 5, and they scattered around the peptide surfaces to enhance structural rigidity. Low-energy structure 68384 had the highest count of both hydrophobic and proline residues, harboring five of each.

Large macrocycle 24 residue designs

Fig 6. Stability analysis of 20-residue designs. (A) Correlation plot between P_{Near} values calculated by Rosetta and our ClusterGen. The ClusterGen's P_{Near} values are plotted against the backbone root-mean-square radii. (B) Example energy landscape comparison. By selecting the lowest-energy structures (marked by a purple cross) as the native states, the ClusterGen landscapes were reshaped. (C) Top designs with minor conformation changes between their initial target states and the lowest-energy structures. (D) Low-energy structures (colored in green) that show major backbone conformation changes from their designed structures (red). These involve formation of a short helix or a compact bending.

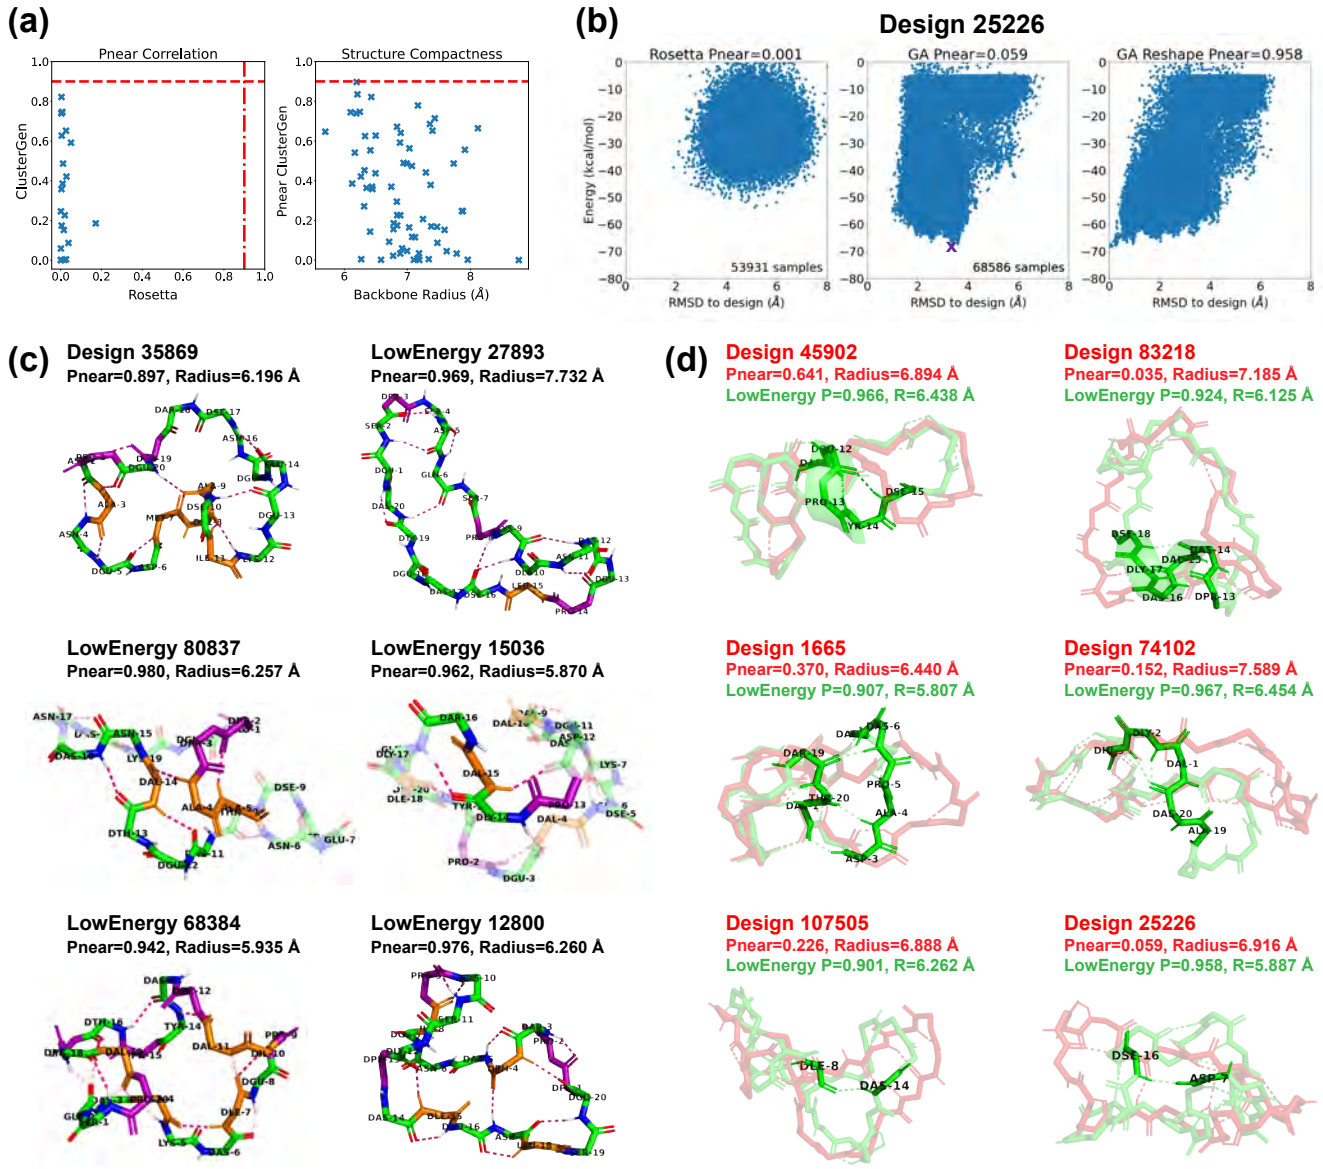

We validated the top 11 designs using both Rosetta and our ClusterGen. Rosetta's failure to thoroughly explore the energy landscapes resulted in near-zero P_{Near} values (Fig 7a). We then validated 71 more designs using ClusterGen, and the highest P_{Near} value obtained is 0.896 from design 31759, with a backbone radius of 7.408 Å. Designs 19384 and 21698 also have high P_{Near} values of 0.790 and 0.786. To expand our search for stable structures, we reshaped the energy landscapes by selecting the lowest-energy samples as the native states and recomputing RMSDs (*e.g.*, Fig 7b). There were 14 low-energy structures achieving $P_{Near} > 0.9$, of which 10 had energies below -75 kcal/mol, including the one for design 31759. We show these 10 structures, as well as the low-energy structures 19384 and 21698

Fig 7. Stability analysis of 24-residue designs. (A) Correlation plot between P_{Near} values calculated by Rosetta and ClusterGen. The ClusterGen's P_{Near} values are plotted against the backbone root-mean-square radii. (B) Example energy landscape comparison. By selecting the lowest energy structures (marked by a purple cross) as the native states, the ClusterGen landscapes were reshaped. (C) Top designs with minor conformation changes in their low energy structures. (D) Low energy structures (colored in green) that show major backbone conformation changes from their designed structures (red).

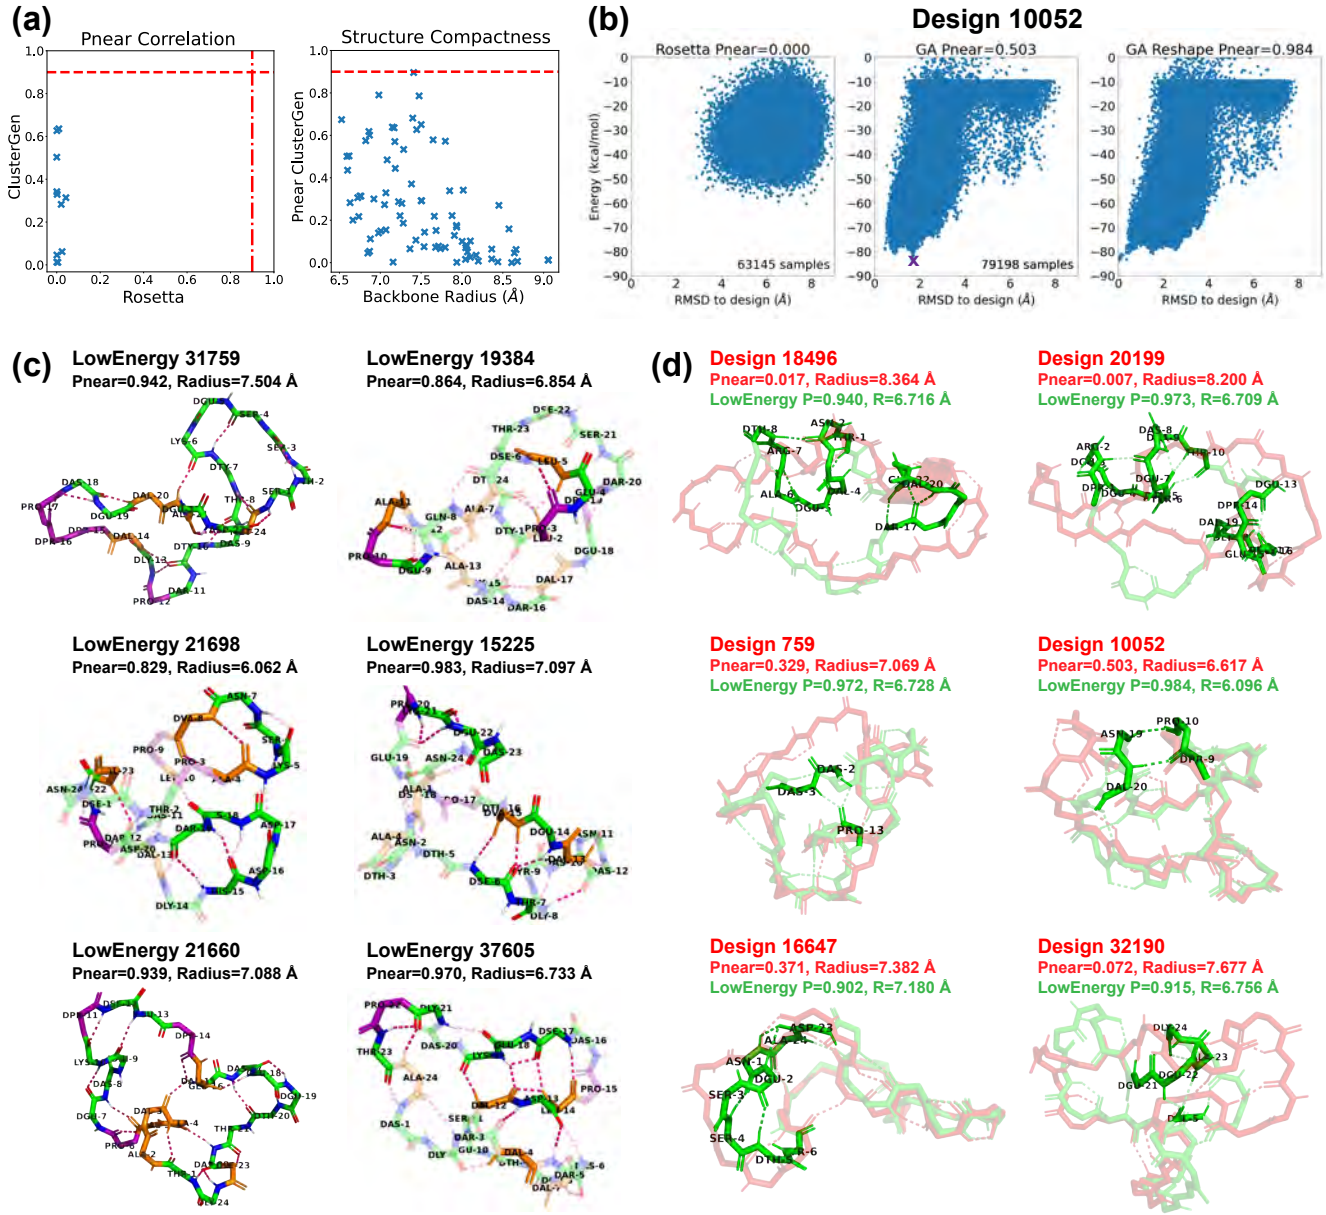

in Fig 7c,d.

Similar to the 20-residue designs, half of these 24-residue designs had only local conformational changes in their low-energy structures (Fig 7c). Structure 21698 showed a diverse array of secondary structures: a 3_{10}-helix in residues 21-23; an isolated β-bridge between residues 4 and 8 with an intermediate backbone turn across residues 5-7; and another β-bridge between residues 15 and 19, enclosing a central backbone turn. The two backbone turn regions were parallel, connected by two

H-bonds, forming a layered configuration. Structure 19384 featured two 3_{10}-helices in residues 3-5 and 9-11. In structure 15225, a β-ladder linked residues 6-7 and 13-15, and consecutive backbone turns formed in residues 20-23. Structure 37605 stood out by forming the highest number of backbone H-bonds (17 in total). Moreover, it had a β-ladder between residues 12-14 and 17-19, combined with an isolated β-bridge between residues 4 and 13. Also, a 3_{10}-helix formed in residues 21-23.

For designs undergoing major conformational shifts from their designed configurations to their low-energy configurations, there were notable decreases in their backbone radii (Fig 7d). The largest radius reductions, approximately 1.5 Å, occured in structures 18496 and 20199. These two low-energy structures adopted shapes resembling a "palm" with three "fingers". The left two "fingers" are held close to each other by an inter-"finger" H-bond between residues 2 and 8 in structure 18496, and residues 2 and 7 in structure 20199. Other designs exhibited more modest decreases in their backbone radii, less than 1 Å. In structures 759, 10052, and 32190, the formation of long-range H-bonds played a crucial role in bridging distant segments, enhancing the structural coherence and stability. Structure 16647 had a 3_{10}-helix forming in the less structured region of residues 2-4.

The top 24-residue structures had a greater proportion of hydrophobic amino acids than the 20-residue ones, featuring between 3 to 8 hydrophobic residues (S8 Fig): in this size range, folds with true hydrophobic cores like those of natural proteins begin to emerge. The number of proline residues remained at the same level of 2-5.

Molecular Dynamics results for top designs

We conducted 1-μs molecular dynamics (MD) simulations on the top designs of 15 residues (Fig 5), 20 residues (Fig 6), and 24 residues (Fig 7), with a timestep of 2 fs. The backbone C^{α}-atom RMSD was measured every 10 ps, comparing the trajectory frame to the initial designed structure. Among the 30 trajectories analyzed, 9 exhibited relatively low RMSDs, indicating kinetic stability. These 9 RMSD trajectories are displayed in Fig 8a and S9 Fig, and the rest are shown in S10 Fig.

The 15-residue design 169032 displayed the most stable trajectory, maintaining an RMSD below 2 Å for the majority of the simulation. Its highest RMSD was 3.51 Å, yet it retained a globally similar shape to the original design. Trajectory 136805 showed oscillating RMSD values between 0.27 and 5.24 Å, with frequent shifts to a slender conformation marked by a bent backbone. Trajectory 3114 presented a symmetrical RMSD profile, with the structure transitioning through two stable phases to adopt an elongated configuration and subsequently reverting to its initial design.

For 20 residues, RMSDs generally fluctuated around the 2-4 Å level. States with low RMSDs preserved structures closely resembling the original designs, while high-RMSD states had more expanded conformations. Trajectory 68384 demonstrated a gradual increase in RMSD, followed by a return to the initial design. Trajectory 45902 exhibited a temporary spike in RMSD.

For 24 residues, trajectory 21698 showed large RMSD oscillations within the first 200 ns. It then stabilized at a 4 Å plateau until around 800 ns, after which the RMSD decreased and the structure reverted to the original design with local conformational changes. Trajectory 19384, similarly, underwent early RMSD fluctuations but levelled off at approximately 4 Å for the rest of the simulation. Trajectory 31759 consistently alternated between a high 4 Å RMSD wide conformation, and a low 2 Å RMSD slender conformation.

Overall, designs 169032, 45902, and 31759 exhibited the most kinetically stable trajectories, showing minor RMSD fluctuations. The remaining designs experienced larger RMSD variations, yet they frequently reverted to the designed structures. To explore the energy landscapes and assess the thermodynamic stability of these designs, we proceeded with replica exchange molecular dynamics (REMD) simulations.

Replica Exchange results for stable trajectories

In addition to the nine designs with stable MD trajectories (Fig 8a), we conducted replica exchange molecular dynamics (REMD) simulations on three PDB crystal structures (8-residue PDB 6ucx, 10-resdiue PDB 6uf7, and 12-residue PDB 6uf8) from a prior study [19] as our positive control group. As

Fig 8. Molecular dynamics simulation results of top designs. (A) Stable 1-μs MD trajectories. RMSDs are calculated between the backbone C^{α} atoms of the trajectory frames and our designed structures. Snapshots are shown for selected time points, with the original designed structure colored in red and the trajectory frame colored in blue. (B) REMD free energy surfaces. From the lowest two free energy basins (marked by green boundaries), representative structures (colored in blue) are extracted from the histogram bins and aligned against our designed structures (red), with the population percentages in the minima labeled aside.

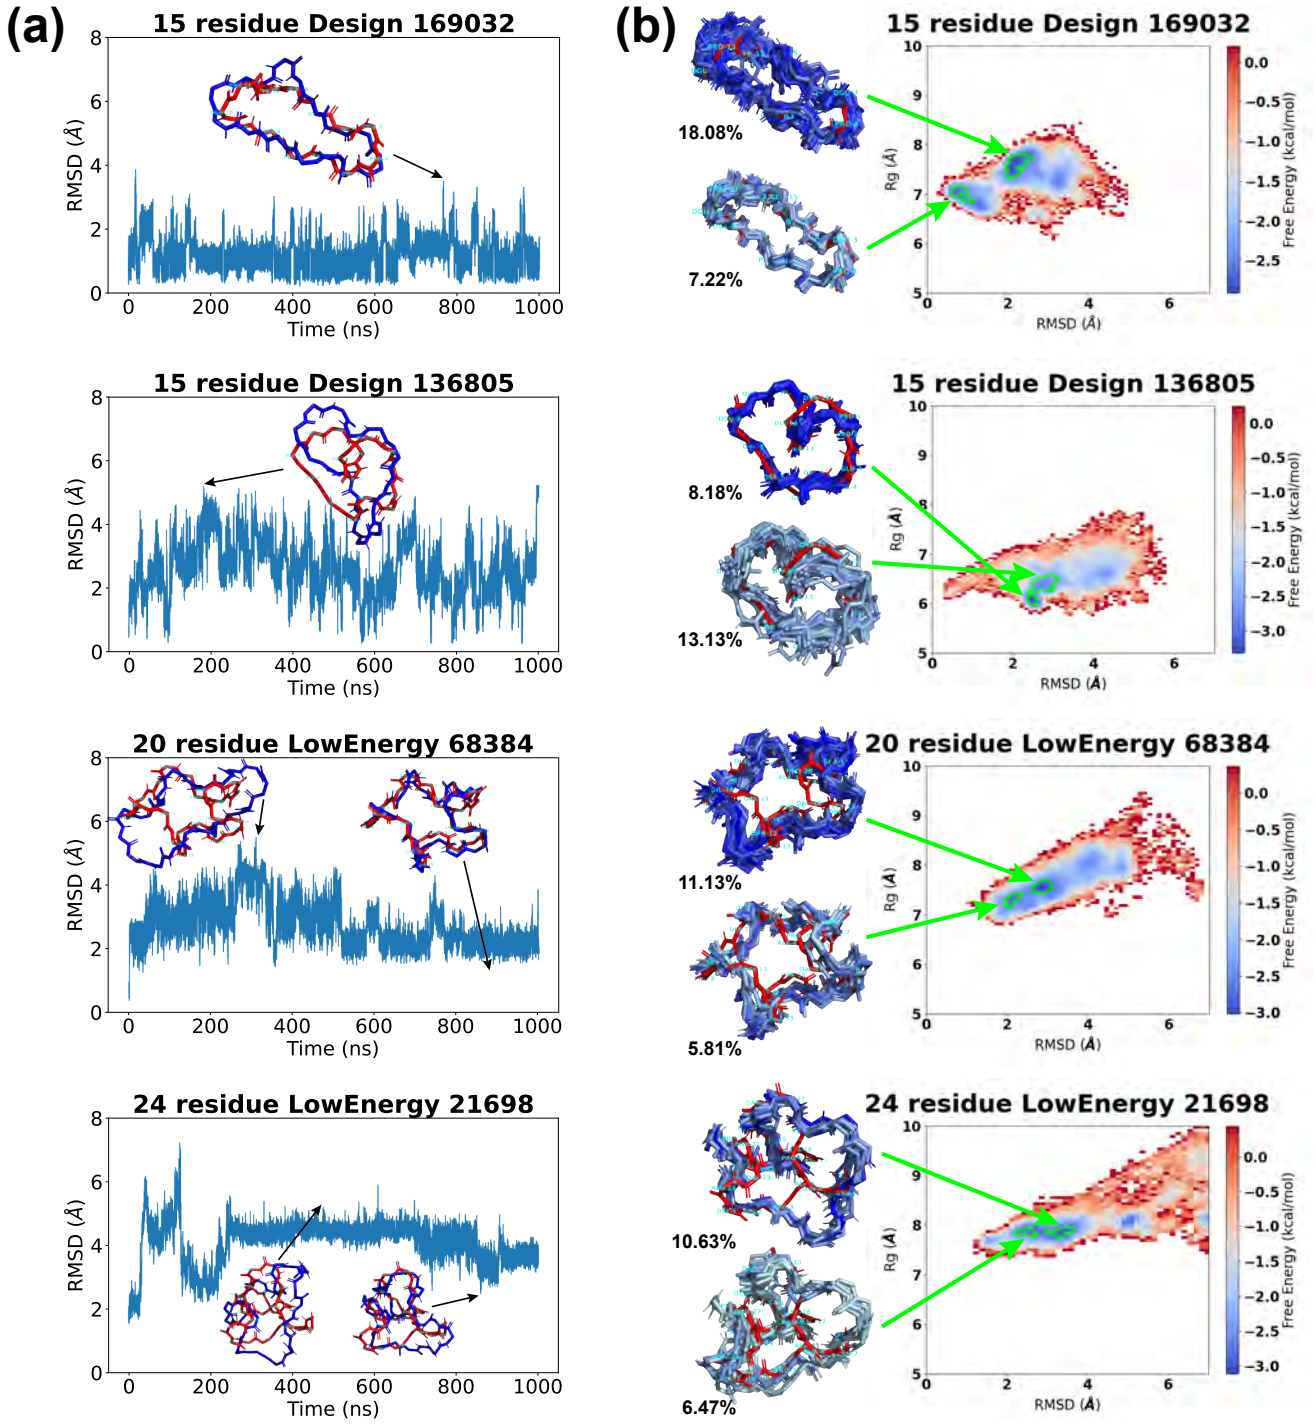

negative controls, we selected one structure from each size of 8, 10, 12, 15, 20, and 24 residues, randomly permuted its amino acid sequences, and refined its side-chains using Rosetta's FastRelax Cartesian relaxation protocol (S9 Appendix). These control groups provided a basis for evaluating our designs' REMD outcomes.

We verified the convergence of the REMD simulations by examining physical properties like temperature and radius of gyration. Details and associated plots are provided in S10 Appendix and S11 Fig-S14 Fig. Subsequently, we calculated the average RMSDs of C^{α}-atoms for uncorrelated configurations sampled at 300 K against the original designs, as documented in Table 1.

Both our positive and negative control groups showed noticeable increases in the average RMSDs as the size grew. For smaller macrocycles (around 15 residues), an RMSD around 2 Å indicates a very strong candidate, around 3 Å is good, 4 Å is weak, and greater than 5 Å is poor, with a greater allowance for larger sizes. Consequently, our 15-residue Design 169032, with an RMSD of just 2.21 Å, stood out as a highly promising candidate. Similarly, the 15-residue Design 136805 and the 20-residue LowEnergy 68384, each with an average RMSD of approximately 3.2 Å, were also favorable. Among the 24-residue designs, LowEnergy 21698, despite an RMSD of 4.04 Å, distinguished itself with a significantly lower RMSD compared to its counterparts, thus making it a viable candidate for our selection.

The free energy surfaces (FES) for these four candidates are shown in Fig 8b. As expected, Design 169032 had energy minima in low-RMSD regions. Its lowest energy minimum E_{min} was located at 2.04 Å RMSD. An 18.08% population of the 300 K configurations fell in this energy basin (energy minimum bin and adjacent bins with free energies $\leq 0.9 * E_{min}$). The second lowest energy basin, at an RMSD of 0.71 Å, accounted for 7.22% of the population. We aligned representative structures from these energy basins (one per bin) with the designed structure, and noted minor deviations in the two proline backbone ends (residues 4 and 11).

Design 136805 exhibited a "heart-like" configuration, with the backbone curling in the middle. Its free energy surface featured a deep concentrated energy basin at 2.49 Å RMSD, accounting for 8.18% of the population. Structures from this energy basin aligned well with our design, showing deviations mainly in the curl region of the backbone (residues 14 and 15). Nearby, the second-lowest energy minimum was found at 2.71 Å RMSD, presenting a broader basin that encompassed 13.13% of the population. Structures from this basin were more variable, with main differences in the right half of the "heart" (residues 1-6).

LowEnergy 68384 displayed a "star-like" conformation with five arms, and its free energy surface exhibited a diagonal distribution. The energy minima were situated in the lower left region, featuring compact structures akin to the original design. The lowest energy basin, at an RMSD of 2.75 Å, held 11.13% of the population. Structures from this basin revealed more expanded conformations in two arms (residues 11-13 and residues 19-20+1-4). The second lowest energy basin, at an RMSD of 2.16 Å,

Table 1. Average RMSDs calculated for REMD simulations. For each simulation, the average RMSD is computed for backbone C^{α} atoms using uncorrelated configurations sampled at temperature state 300 K.

n	Design name	Average RMSD (Å)	Random RMSD (Å)
8	PDB 6ucx	0.48	2.23
10	PDB 6uf7	1.31	2.92
12	PDB 6uf8	2.77	3.01
15	**Design 169032**	**2.21**	4.49
15	**Design 136805**	**3.24**	
15	Design 3114	3.63	
20	**LowEnergy 68384**	**3.23**	4.94
20	LowEnergy 83218	4.04	
20	LowEnergy 45902	4.78	
24	**LowEnergy 21698**	**4.04**	
24	LowEnergy 19384	5.10	
24	LowEnergy 31759	5.67	5.94

comprised 5.81% of the population and included structures that closely matched the overall shape of our design.

LowEnergy 21698 featured a narrow low-energy region spanning RMSDs of 2 to 4 Å. In it, the lowest energy basin, at an RMSD of 3.23 Å, accounted for 10.63% of the population. The corresponding structures exhibit a structural opening in residues 3-10. Adjacent to this, the second-lowest energy basin at RMSD 2.34 Å held 6.47% of the population and displayed local structural deviations in residues 7-9 compared to the original design.

For other designs with high average RMSDs, their free energy surfaces generally showed minima in regions of high RMSD, as depicted in S15 Fig. However, the 20-residue LowEnergy 83218 and 45902 exhibited disperse FES, with the two lowest energy basins located at RMSDs of approximately 2 Å and 5 Å, totaling around 7% of the population. Regarding the random sequences, there were notable changes in the FES compared to the designed sequences, particularly with the 8-residue 6ucx, which displayed multiple energy minima in its random FES, while a single minimum in its design FES. This suggests the importance of sequence design in maintaining a thermodynamically stable fold.

Discussion

In this work, we have introduced a pipeline, called CyclicChamp, for cyclic peptide design (Fig 1). Many past works have used a single-shot generalized kinematic closure (GenKIC) algorithm to sample closed macrocycle conformations [4, 5, 10, 15, 19, 20, 34, 35]. Unfortunately, the GenKIC technique limits the size of macrocycle for which the conformation space may be extensively explored, either for design or for validation.

By contrast, CyclicChamp performs an iterative search of cyclic backbones with favorable features like strong H-bonds and no steric clashes. Because this produces more viable backbone conformations in less time, CyclicChamp is able to design small macrocycles at lower computational cost, and for the first time, access sizes as large as 24 residues without relying on symmetry or chemical cross-links to limit the accessible conformational space. The basic insight is to transform the cyclic backbone constraint into an error function to allow the use of optimization methods like simulated annealing and genetic algorithms. The optimal simulated annealing parameters were selected from well-spaced random samples of possible parameter value combinations, obtained using combinatorial design [27]. While we have assumed ideal bond angles, bond lengths, and ω torsion angles to simplify the evaluation of the cyclic error function (Fig 1), generalizations that allow these degrees of freedom to deviate slightly from ideal values during the backbone simulated annealing time steps are possible.

Using these algorithmic ideas, we have generated macrocycles of four sizes. For 7-residue designs, we conducted a comprehensive search of Ramachandran spaces by considering all possible torsion bin center combinations for initial backbone torsion angles. Because the number of torsion bin combinations grows exponentially, for larger designs of 15-24 residues, we randomly selected 100,000 initial combinations. Large pools of backbone candidates with distinct torsion bin strings were generated for 15 and 20 residues (Fig 2).

The sparse clusters found in the 7-residue design were due to the limited torsion bin strings. For 24 residues, the exponential scaling of the accessible conformational space with number of backbone degrees of freedom, and the fact that only a tiny portion of the conformation space represents backbones with favourable features (e.g. hydrogen bonds) that could be stabilized by suitable choice of sequence, meant that our backbone simulated annealing algorithm reached its limit. Because the search space grows exponentially, and because solving even simpler discrete version of such problem is NP-complete [36], there is no known efficient means of sampling conformations across all sizes, though better heuristic methods like ours can increase the maximum size of peptide that can be sampled and designed.

Future studies might experiment with alternative energy models or consider less stringent requirements for cyclic error and repulsive energy when selecting backbone candidates to try to push this limit higher.

After the relaxation and design steps were applied on the clustered backbones, we conducted stability tests on the designs having the lowest energies. For 7-residue designs, both Rosetta's random sampling

method and our Ramachandran-stability filtering method were employed to generate energy landscapes. A positive correlation was found to exist between the P_{Near} values computed by the two methods. We observed instances in which the filtering method explored the low RMSD regions in the energy landscapes more thoroughly than the random sampling method.

We noted that the optimal energy range for high P_{Near} values did not always correspond to the lowest energy levels (Fig 3). Design calculations necessarily consider only the desired conformation that one is stabilizing, in order to make the problem tractable; however, the true problem that one wishes to solve is that of maximizing the energy gap between the desired conformation and all alternative conformations. The lack of correlation between the best P_{Near} values and the lowest single-state energies could be that Rosetta has artificial ways of lowering the energy of the designed state, such as adding hydrophobic groups, which tend to stabilize all structures universally instead of uniquely stabilizing the designed structure and maximizing the energy gap between this and alternative states.

The Ramachandran-stability filtering approach can extend to design cyclic peptides with constrained but not fully specified sequences. For instance, to design a stable 7-residue cyclic peptide with alanine residues as the first and fifth amino acids, from the backbone candidate pool sampled by layered simulated annealing, we can identify backbones whose first and fifth residues' torsion angles fall in the Ramachandran space accessible to alanine. This approach allows us to reuse existing pools of backbone candidates.

Starting from 15 residues, Rosetta's method tended to struggle with exploring the low-RMSD regions, and often generated similar round-shaped energy landscapes (Fig 5). To resolve this issue, our method's ClusterGen algorithm begins with two simulated annealing runs targeting low energy and low RMSD, effectively broadening the RMSD spectrum of the landscape. The subsequent genetic algorithm identifies energy minima through iterations of crossover, mutation, and selection. ClusterGen has successfully differentiated designs of various energy landscape shapes (Fig 5).

In the top 15-residue designs, we observed short alpha helices as depicted in Fig 5. Recurring backbone bendings were induced by $i, i+3$ H-bonds, which led to more twisted shapes compared to the simple circular backbones seen in 7-residue designs. In the top 20- and 24-residue designs, we saw more diverse secondary structures such as 3_{10}-helices, β-bridges, and β-ladders (Fig 6, Fig 7). Although complete α-helices or β-sheets are not fully formed, fragments of these structures start appearing, aiding in the stabilization of these mid-sized peptides. Long-range H-bonds also play a crucial role in stabilizing the 20 and 24 residue macrocycles.

Additionally, for 15-24 residues, we see an enrichment of high P_{Near} values in compact structures (Fig 5, Fig 6, Fig 7). This suggests that simply sorting the designs in ascending order of energies may not be the most effective strategy to identify top designs for stability validation. A more nuanced approach could be to select designs that feature secondary structures or long-range H-bonds and have backbone radii below a specific threshold, and then to sort them by energy.

We have found close backbone matches in our 7-residue designs to the three previously-published, experimentally-solved Rosetta designs (Fig 4). For 15, 20, and 24 residues, we conducted MD and REMD simulations to evaluate the kinetic and thermodynamic stability of our top designs (Fig 8). Specifically, the 15-residue Design 169032 demonstrated exceptional stability, with free energy minima around the 2 Å RMSD region, and the representative structures closely aligning with the design. Another 15-residue Design 136805 and a 20-residue LowEnergy 68384 managed to preserve their overall shapes throughout the simulation, despite some local conformational movements. The absence of a robust 24-residue candidate underscores the escalating challenge in design complexity. While expanding the pool of initial points for backbone sampling might help to find more stable designs at some sizes, the exponentially growing search spaces will inevitably render random conformational sampling intrinsically unproductive for very large peptides. Reintroduction of past strategies (symmetry, disulfides, other cross-links) to limit accessible degrees of freedom will still likely be needed for very large peptides.

Nevertheless, to the best of our knowledge, this work represents the first instance of general, unconstrained design of 15-, 20-, and 24-residue macrocycles without relying on limitation of degrees of freedom through the use of symmetry, disulfides, or other cross-links. The capability to design such large sizes not only enhances the structural diversity of cyclic peptides for future drug search, but also allows

larger interaction surfaces for drug binding. Moreover, this opens a door to the design of cyclic-peptide enzymes, which require larger sizes to form active site pockets and may incorporate exotic chemical groups with active-site residues for catalysis.

Supporting information

S1 Appendix. Cyclic error derivation.

S2 Appendix. Backbone energy functions.

S3 Appendix. Layered simulated annealing.

S4 Appendix. Combinatorial design.

S5 Appendix. Torsion angle FastRelax script.

S6 Appendix. Small macrocycle FastDesign script.

S7 Appendix. Large macrocycle FastDesign script.

S8 Appendix. ClusterGen stability analysis.

S9 Appendix. Cartesian coordinate FastRelax script.

S10 Appendix. MD and REMD protocol.

S1 Fig. Backbone energy functions.

S2 Fig. Extended glycine Ramachandran space sampling.

S3 Fig. Ramachandran spaces for L and D amino acids.

S4 Fig. ClusterGen crossover and mutation.

S5 Fig. ClusterGen computation time breakdown.

S6 Fig. CyclicChamp design results.

S7 Fig. Top 20-residue designs shown in sphere mode.

S8 Fig. Top 24-residue designs shown in sphere mode.

S9 Fig. Molecular dynamics simulation stable trajectories.

S10 Fig. Molecular dynamics simulation unstable trajectories.

S11 Fig. REMD temperature plots.

S12 Fig. REMD temperature dwell time plots.

S13 Fig. REMD convergence check of Rg.

S14 Fig. REMD convergence check of RMSD.

S15 Fig. REMD free energy surfaces.

S1 Table Backbone simulated annealing parameters.

Acknowledgments

We extend our sincerest gratitude to several individuals whose expertise help the successful completion of this work. We would like to thank P. Douglas Renfrew at the Flatiron institute for his suggestions on our data analysis, Bargeen Turzo for the MD package setup, Pilar Cossio, Sonya Hanson, Justin Lindsay, and Miro Astore for their guidance on MD and REMD simulations, and Nick Carriero and Géraud Krawezik for enhancing our codes' parallelism on the Flatiron Institute's HPC. We also would like to thank Shenglong Wang for setting up the NYU HPC. The Flatiron Institute is a division of the Simons Foundation. QZ and VKM were funded by the Simons Foundation. DS was funded by the US National Science Foundation grants 1840761, 2304758, 1934388, the National Institutes of Health grant 1R01GM121753-01A1, and NYU Wireless.

References

1. Zorzi A, Deyle K, Heinis C. Cyclic peptide therapeutics: past, present and future. Curr Opin Chem Biol. 2017;38:24–29.
2. Mulligan V. The emerging role of computational design in peptide macrocycle drug discovery. Expert Opin Drug Discov. 2020;15:833–852.
3. Joo S. Cyclic peptides as therapeutic agents and biochemical tools. Biomol Ther. 2012;20:19–26.
4. Mulligan VK, Workman S, Sun T, Rettie S, Li X, Worrall LJ, et al. Computationally designed peptide macrocycle inhibitors of New Delhi metallo-β-lactamase 1. Proc Natl Acad Sci USA. 2021;118:e2012800118.
5. Hosseinzadeh P, Watson PR, Craven TW, Li X, Rettie S, Pardo-Avila F, et al. Anchor extension: a structure-guided approach to design cyclic peptides targeting enzyme active sites. Nat Commun. 2021;12:3384.
6. D'Aloisio V, Dognini P, Hutcheon G, Coxon C. PepTherDia: database and structural composition analysis of approved peptide therapeutics and diagnostics. Drug Discov Today. 2021;26:1409–1419.
7. Rettie S, Campbell K, Bera A, Kang A, Kozlov S, et al. Cyclic peptide structure prediction and design using AlphaFold. 2023;doi:bioRxiv:10.1101/2023.02.25.529956.
8. Baek M, Dimaio F, Anishchenko I, Dauparas J, et al. Accurate prediction of protein structures and interactions using a three-track neural network. Science. 2021;373:871–876.
9. Tsaban T, Varga JK, Avraham O, Ben-Aharon Z, Khramushin A, Schueler-Furman O. Harnessing protein folding neural networks for peptide–protein docking. Nat Commun. 2022;13:176.
10. Hosseinzadeh P, Bhardwaj G, Mulligan V, Shortridge M, Craven T, Pardo-Avila F, et al. Comprehensive computational design of ordered peptide macrocycles. Science. 2017;358:1461–1466.

11. Georgiev I, Lilien R, Donald B. The minimized dead-end elimination criterion and its application to protein redesign in a hybrid scoring and search algorithm for computing partition functions over molecular ensembles. J Comput Chem. 2008;29:1527–1542.

12. Kaur G, Kapoor S, Kaundal S, Dutta D, Thakur KG. Structure-guided designing and evaluation of peptides targeting bacterial transcription. Front Bioeng Biotechnol. 2020;8:797.

13. Mulligan V. Current directions in combining simulation-based macromolecular modeling approaches with deep learning. Expert Opin Drug Discov. 2021;16:1025–1044.

14. Mandell DJ, Coutsias EA, Kortemme T. Sub-angstrom accuracy in protein loop reconstruction by robotics-inspired conformational sampling. Nat Methods. 2009;6:551–552.

15. Bhardwaj G, Mulligan V, Bahl C, Gilmore J, Harvey P, Cheneval O, et al. Accurate de novo design of hyperstable constrained peptides. Nature. 2016;538:329–335.

16. Alford R, Leaver-Fay A, Jeliazkov J, O'Meara M, et al. The Rosetta all-atom energy function for macromolecular modeling and design. J Chem Theory Comput. 2017;13:3031–3048.

17. Leaver-Fay A, Tyka M, Lewis SM, Lange OF, Thompson J, Jacak R, et al. ROSETTA3: an object-oriented software suite for the simulation and design of macromolecules. 2011;487:545–574. doi:10.1016/B978-0-12-381270-4.00019-6.

18. Mulligan VK, Hosseinzadeh P. 3. In: Computational design of peptide-based binders to therapeutic targets;. p. 55–102. Available from: `https://pubs.acs.org/doi/abs/10.1021/bk-2022-1417.ch003`.

19. Mulligan V, Kang C, Sawaya M, Rettie S, Li X, Antselovich I, et al. Computational design of mixed chirality peptide macrocycles with internal symmetry. Protein Sci. 2020;29:2433–2445.

20. Dang B, Wu H, Mulligan VK, Mravic M, Wu Y, Lemmin T, et al. De novo design of covalently constrained mesosize protein scaffolds with unique tertiary structures. Proc Natl Acad Sci USA. 2017;114:10852–10857.

21. Mulligan V. Current directions in combining simulation-based macromolecular modeling approaches with deep learning. Expert Opin Drug Discov. 2021;16:1025–1044.

22. Khatib F, Cooper S, Tyka M, Xu K, et al. Algorithm discovery by protein folding game players. Proc Natl Acad Sci USA. 2011;108:18949–18953.

23. Maguire J, Haddox H, Strickland D, Halabiya S, Coventry B, Griffin J, et al. Perturbing the energy landscape for improved packing during computational protein design. Proteins: Struct, Funct, Bioinf. 2021;89:436–449.

24. Coutsias E, Seok C, Wester M, Dill K. Resultants and loop closure. Int J Quantum Chem. 2006;106:176–189.

25. Go N, Scheraga H. Ring closure and local conformational deformations of chain molecules. Macromolecules. 1969;3:178–187.

26. Burnside W. On Some Properties of Groups of Odd Order. Proc Natl Acad Sci USA. 1900;33:162–184.

27. Colbourn C, Martirosyan S, Mullen G, Shasha D, Sherwood G, Yucas J. Products of mixed covering arrays of strength two. J Combin Designs. 2006;14:124–138.

28. Kabsch W. A solution for the best rotation to relate two sets of vectors. Acta Cryst. 1976;32:922–923.

29. Eastman P, Swails J, Chodera J, McGibbon R, Zhao Y, et al. OpenMM 7: Rapid development of high performance algorithms for molecular dynamics. PLOS Comp Biol. 2017;13:e1005659.

30. Sugita Y, Okamoto Y. Replica-exchange molecular dynamics method for protein folding. Chem Phys Lett. 1999;314:141–151.

31. Chodera J, Rizzi A, Naden L, Beauchamp K, Grinaway P, et al. Choderalab/openmmtools: 0.23.1; 2023.

32. McGibbon R, Beauchamp K, Harrigan M, Klein C, Swails J, et al. MDTraj: A modern open library for the analysis of molecular dynamics trajectories. Biophys J. 2015;109:1528–1532.

33. Harris C, Millman K, Walt S, Gommers R, Virtanen P, et al. Array programming with NumPy. Nature. 2020;585:357–362.

34. Chorsi MS, Linthicum W, Pozhidaeva A, Mundrane C, Mulligan VK, Chen Y, et al. Ultra-confined controllable cyclic peptides as supramolecular biomaterials. Nano Today. 2024;56:102247. doi:10.1016/j.nantod.2024.102247.

35. Bhardwaj G, O'Connor J, Rettie S, Huang YH, Ramelot TA, Mulligan VK, et al. Accurate de novo design of membrane-traversing macrocycles. Cell. 2022;185(19):3520–3532.e26. doi:10.1016/j.cell.2022.07.019.

36. Pierce NA, Winfree E. Protein Design is NP-hard. Protein Eng Des Sel. 2002;15:779–782.